# HP SureStore E Disk Array FC60
# Advanced User's Guide

This manual was downloaded from http://www.hp.com/support/fc60/




Edition E1200
Printed in U.S.A.

## Notice

© Hewlett-Packard Company, 1999, 2000. All rights reserved.

Hewlett-Packard Company makes no warranty of any kind with regard to this document, including, but not limited to, the implied warranties of merchantability and fitness for a particular purpose. Hewlett-Packard shall not be liable for errors contained herein or for incidental or consequential damages in connection with the furnishing, performance, or use of this material.

This document contains proprietary information, which is protected by copyright. No part of this document may be photocopied, reproduced, or translated into another language without the prior written consent of Hewlett-Packard. The information contained in this document is subject to change without notice.

## Trademark Information

Microsoft, Windows, Windows NT, and Windows 2000 are registered trademarks of the Microsoft Corporation.

## Safety Notices




To protect against personal injury and product damage, do not attempt to lift the product without the assistance of another person or lift device.



Components bearing this symbol may be hot to touch.



Components bearing this symbol are fragile. Handle with care.



Components bearing this symbol are susceptible to damage by static electricity. ESD precautions are required.

### Service

Any servicing, adjustment, maintenance, or repair must be performed only by authorized service-trained personnel.

## Format Conventions

| | Denotes |
|---|---|
| **WARNING** | **A hazard that can cause personal injury** |
| ***Caution*** | A hazard that can cause hardware or software damage |
| **Note** | Significant concepts or operating instructions |
| **this font** | Text to be typed verbatim: all commands, path names, file names, and directory names |
| `this font` | Text displayed on the screen |

## Printing History

1st Edition - September 1999
2nd Edition - October 1999
3rd Edition - February 2000
4th Edition - July 2000
5th Edition - September 2000
6th Edition - October 2000
7th Edition - December 2000

## Manual Revision History

### December 2000

| Change | Page |
|---|---|
| Added Figure 87 to clarify operation of the write cache flush thresholds. | 253 |
| Added note regarding the impact of LUN binding on performance. | 250 |
| Added information on Managing the Universal Transport Mechanism (UTM). | 298 |
| Added information on major event logging available with firmware HP08. | 307 |
| Added Allocating Space for Disk Array Logs section describing use of environment variable `AM60_MAX_LOG_SIZE_MB.` | 308 |
| Added information on Purging Controller Logs. | 311 |
| Added information for RAID 0 support on HP-UX. | 47 |
| Changed the required minimum number of disk modules per enclosure from 2 to 4 based on power supply requirements for the disk enclosure. | 73 |

## About This Book

This guide is intended for use by system administrators and others involved in operating and managing the HP SureStore E Disk Array FC60. It is organized into the following chapters and section.

| | |
|---|---|
| **Chapter 1, Product Description** | Describes the features, controls, and operation of the disk array. |
| **Chapter 2, Topology and Array Planning** | Guidelines for designing the disk array configuration that best meets your needs. |
| **Chapter 3, Installation** | Instruction for moving the disk array. |
| **Chapter 4, Managing the Disk Array on HP-UX** | Complete instructions for managing your disk array using the available management software. |
| **Chapter 5, HP-UX Diagnostic Tools** | Information on using STM to gather information about disk array status. |
| **Chapter 6, Troubleshooting** | Instructions for isolating and solving common problems that may occur during disk array operation. |
| **Chapter 7, Removal and Replacement** | Instructions for removing and replacing all customer-replaceable components. |
| **Chapter 8, Reference / Legal / Regulatory** | Regulatory, environmental, and other reference information. |
| **GLOSSARY** | |
| **Index** | |

## Related Documents and Information

The following items contain information related to the installation and use of the HP SureStore E Disk Array and its management software.

- *HP SureStore E Disk Array FC60 Advanced User's Guide* - this is the expanded version of the book you are reading. Topics that are discussed in more detail in the Advanced User's Guide are clearly identified throughout this book.
   Download: www.hp.com/support/fc60
- *HP Storage Manager 60 User's Guide* - this guide describes the features and operation of the disk array management software for Windows NT and Windows 2000. It is included with the A5628A software kit.
   Download: www.hp.com/support/fc60
- *HP Storage Manager 60 Introduction Guide* - this guide introduces the disk array management software for Windows NT and Windows 2000. It is included in electronic format on the *HP Storage Manager 60 CD*.
   Download: www.hp.com/support/fc60
- *Fibre Channel Mass Storage Adapters Service and User Manual* (A3636-90002) - describes Fibre Channel Mass Storage/9000. It describes installation of the Fibre Channel I/O adaptors into K-, D-, T-, and V-class systems.
   Download: www.hp.com/essd/efc/A3636A_documentation.html
- *Using EMS HA Monitors* (B5735-90001) - contains information about the EMS environment used for hardware monitoring.
   Download: http://www.docs.hp.com/hpux/ha/
- *EMS Hardware Monitors User's Guide* - describes how to use the EMS Hardware Monitors to protect your system from undetected failures.
   Download: http://www.docs.hp.com/hpux/systems/
- *Diagnostic/IPR Media User's Guide* (B6191-90015) - provides information on using STM, and enabling the EMS Hardware Event Monitors.
   Download: http://www.docs.hp.com/hpux/systems/
- *Managing MC/ServiceGuard* (B3939-90024) - provides information on creating package dependencies for hardware resources.
   Download: http://www.docs.hp.com/hpux/ha/

# 1 Product Description


Product Description . . . . . 18
Operating System Support . . . . . 20
Management Tools . . . . . 20
Features . . . . . 20
High Availability . . . . . 21
Scalable Storage Capacity . . . . . 21
LED Status Monitoring . . . . . 22
EMS Hardware Event Monitoring (HP-UX Only) . . . . . 22
Disk Enclosure Components . . . . . 23
Operation Features . . . . . 25
Disk Enclosure SC10 Modules . . . . . 27
Array Controller Enclosure Components . . . . . 34
Front Cover . . . . . 37
Controller Modules . . . . . 38
Controller Fan Modules . . . . . 40
Power Supply Modules . . . . . 42
Power Supply Fan Module . . . . . 43
Battery Backup Unit . . . . . 45
Disk Array High Availability Features . . . . . 47
RAID Technology . . . . . 47
Disk Mirroring . . . . . 47
Data Parity . . . . . 48
Data Striping . . . . . 49
RAID Levels . . . . . 50
RAID Level Comparisons . . . . . 57
Global Hot Spare Disks . . . . . 61
Primary and Alternate I/O Paths . . . . . 64
Capacity Management Features . . . . . 65
LUNs . . . . . 65

Disk Groups . . . . . 65
Disk Array Caching . . . . . 66
Dynamic Capacity Expansion . . . . . 67

# 2 Topology and Array Planning

Overview . . . . . 70
Array Design Considerations . . . . . 71
Array Hardware Configuration . . . . . 71
RAID, LUNs, and Global Hot Spares . . . . . 72
High Availability . . . . . 72
Performance . . . . . 73
Storage Capacity . . . . . 75
Expanding Storage Capacity . . . . . 75
Recommended Disk Array Configurations . . . . . 77
Configuration Considerations . . . . . 77
One Disk Enclosure Configuration . . . . . 78
Two Disk Enclosure Configurations . . . . . 80
Three Disk Enclosure Configurations . . . . . 84
Four Disk Enclosure Configurations . . . . . 88
Five Disk Enclosure Configurations . . . . . 92
Six Disk Enclosure Configurations . . . . . 96
Total Disk Array Capacity . . . . . 100
For high-availability, one disk per SCSI channel is used as a global hot spare. . . 101
Topologies for HP-UX . . . . . 102
Basic Topology . . . . . 103
Single-System Distance Topology . . . . . 110
High Availability Topology . . . . . 115
High Availability, Distance, and Capacity Topology . . . . . 120
Campus Topology . . . . . 125
Performance Topology with Switches . . . . . 129
Topologies for Windows NT and Windows 2000 . . . . . 131

Non-High Availability Topologies . . . . . 133
High Availability Topologies . . . . . 137

# 3 Installation

Overview . . . . . 144
Host System Requirements . . . . . 145
HP-UX . . . . . 145
Windows NT and Windows 2000 . . . . . 146
Site Requirements . . . . . 147
Environmental Requirements . . . . . 147
Electrical Requirements . . . . . 147
Power Distribution Units (PDU/PDRU) . . . . . 150
Installing PDUs . . . . . 152
Recommended UPS Models . . . . . 152
Installing the Disk Array FC60 . . . . . 155
Installing the Disk Enclosures . . . . . 160
Step 1: Collect Required Tools . . . . . 160
Step 2: Unpack the Product . . . . . 160
Step 3: Install Mounting Rails . . . . . 162
Step 4: Install the Disk Enclosure . . . . . 162
Step 5: Install Disks and Fillers . . . . . 166
Moving a Disk Enclosure from One Disk Array to Another . . . . . 168
Installing the Controller . . . . . 170
Step 1: Gather Required Tools . . . . . 170
Step 2: Unpack the Product . . . . . 170
Step 3: Install Mounting Rails . . . . . 173
Step 4: Install the Controller Enclosure . . . . . 173
Configuration Switches . . . . . 176
Disk Enclosure (Tray) ID Switch . . . . . 176
Disk Enclosure DIP Switch . . . . . 176
Fibre Channel Host ID Address Setting . . . . . 179

Attaching Power Cords . . . . . 183
Attaching SCSI Cables and
Configuring the Disk Enclosure Switches . . . . . 187
Full-Bus Cabling and Switch Configuration . . . . . 188
Split-Bus Switch and Cabling Configurations . . . . . 191
Connecting the Fibre Channel Cables . . . . . 196
Applying Power to the Disk Array . . . . . 198
Verifying Disk Array Connection . . . . . 206
On Windows NT and Windows 2000 . . . . . 206
On HP-UX . . . . . 206
Interpreting Hardware Paths . . . . . 208
Installing the Disk Array FC60 Software (HP-UX Only) . . . . . 213
System Requirements . . . . . 213
Verifying the Operating System . . . . . 214
Installing the Disk Array FC60 Software . . . . . 214
Downgrading the Disk Array Firmware for HP-UX 11.11 Hosts . . . . . 215
Configuring the Disk Array . . . . . 216
HP-UX . . . . . 216
Windows NT and Windows 2000 . . . . . 220
Using the Disk Array FC60 as a Boot Device (HP-UX Only) . . . . . 222
Solving Common Installation Problems . . . . . 223
Adding Disk Enclosures to Increase Capacity . . . . . 224
General Rules for Adding Disk Enclosures to the Disk Array . . . . . 224
Step 1. Plan the Expanded Configuration . . . . . 225
Step 2. Backup All Disk Array Data . . . . . 226
Step 3. Prepare the Disk Array for Shut Down . . . . . 226
Step 4. Add the New Disk Enclosures . . . . . 227
Step 5. Completing the Expansion . . . . . 230
Capacity Expansion Example . . . . . 232

# 4 Managing the Disk Array on HP-UX


Tools for Managing the Disk Array FC60 . . . . . 238
System Administration Manager (SAM) . . . . . 238
Array Manager 60 . . . . . 238
STM . . . . . 238
Installing the Array Manager 60 Software . . . . . 240
AM60Srvr Daemon . . . . . 241
Running Array Manager 60 . . . . . 241
Managing Disk Array Capacity . . . . . 242
Configuring LUNs . . . . . 242
Selecting Disks for a LUN . . . . . 243
Assigning LUN Ownership . . . . . 247
Selecting a RAID Level . . . . . 247
Global Hot Spares . . . . . 248
Setting Stripe Segment Size . . . . . 249
Evaluating Performance Impact . . . . . 250
Adding Capacity to the Disk Array . . . . . 254
Adding More Disk Modules . . . . . 254
Adding Additional Disk Enclosures . . . . . 256
Replacing Disk Modules with Higher Capacity Modules . . . . . 256
Upgrading Controller Cache to 512 Mbytes . . . . . 258
Managing the Disk Array Using SAM . . . . . 260
Checking Disk Array Status . . . . . 261
Assigning an Alias to the Disk Array . . . . . 264
Locating Disk Modules . . . . . 265
Binding a LUN . . . . . 267
Unbinding a LUN . . . . . 271
Replacing a LUN . . . . . 271
Adding a Global Hot Spare . . . . . 273
Removing a Global Hot Spare . . . . . 275

Managing the Disk Array Using Array Manager 60 . . . . . 276
Command Syntax Conventions . . . . . 279
Array Manager 60 man pages . . . . . 279
Quick Help . . . . . 279
Selecting a Disk Array and Its Components . . . . . 280
Preparing to Manage the Disk Array . . . . . 281
Checking Disk Array Status . . . . . 282
Managing LUNs . . . . . 289
Managing Global Hot Spares . . . . . 295
Managing Disk Array Configuration . . . . . 297
Managing the Universal Transport Mechanism (UTM) . . . . . 298
Managing Cache Parameters . . . . . 299
Performing Disk Array Maintenance . . . . . 304
Managing Disk Array Logs . . . . . 307
Upgrading Disk Firmware . . . . . 313
Managing the Disk Array Using STM . . . . . 314
Checking Disk Array Status Information . . . . . 314
Binding a LUN . . . . . 314
Unbinding a LUN . . . . . 315
Adding a Global Hot Spare . . . . . 315
Removing a Global Hot Spare . . . . . 315
Locating Disk Modules . . . . . 316
Status Conditions and Sense Code Information . . . . . 317
LUN Status Conditions . . . . . 317
Disk Status Conditions . . . . . 319
Component Status Conditions . . . . . 321
FRU Codes . . . . . 322
SCSI Sense Codes . . . . . 327

## 5 HP-UX Diagnostic Tools

Overview . . . . . 346

Support Tools Manager . . . . . 347
STM User Interfaces . . . . . 347
STM Tools . . . . . 351
Using the STM Information Tool . . . . . 352
Interpreting the Information Tool Information Log . . . . . 354
Interpreting the Information Tool Activity Log . . . . . 354
Using the STM Expert Tool . . . . . 355

# 6 Troubleshooting

Introduction . . . . . 360
About Field Replaceable Units (FRUs) . . . . . 361
HP-UX Troubleshooting Tools . . . . . 362
Windows NT and Windows 2000 Troubleshooting Tools . . . . . 362
EMS Hardware Event Monitoring (HP-UX Only) . . . . . 362
Disk Array Installation/Troubleshooting Checklist . . . . . 365
Power-Up Troubleshooting . . . . . 366
Controller Enclosure Troubleshooting . . . . . 367
Introduction . . . . . 367
Controller Enclosure LEDs . . . . . 368
Master Troubleshooting Table . . . . . 370
SureStore E Disk System SC10 Troubleshooting . . . . . 376
Disk Enclosure LEDs . . . . . 376
Losing LUN 0 . . . . . 376
Interpreting Component Status Values (HP-UX Only) . . . . . 379
Isolating Causes . . . . . 380

# 7 Removal and Replacement

Overview . . . . . 384
Disk Enclosure Modules . . . . . 386
Disk Module or Filler Module . . . . . 386

Disk Enclosure Fan Module . . . . . 392
Disk Enclosure Power Supply Module . . . . . 394
Controller Enclosure Modules . . . . . 396
Front Cover Removal/Replacement . . . . . 397
Controller Fan Module . . . . . 398
Battery Backup Unit (BBU) Removal/Replacement. . . . . 400
Power Supply Fan Module Removal/Replacement. . . . . 403
Power Supply Module Removal/Replacement . . . . . 405
SCSI Cables . . . . . 407

# 8 Reference / Legal / Regulatory

System Requirements. . . . . 410
Host Systems . . . . . 410
Supported Operating Systems. . . . . 410
Fibre Channel Host Adapters . . . . . 411
Models and Options . . . . . 412
A5277A/AZ Controller Enclosure Models and Options . . . . . 412
A5294A/AZ Disk Enclosure SC10 Models and Options . . . . . 414
Disk Array FC60 Upgrade and Add-On Products . . . . . 416
PDU/PDRU Products . . . . . 417
Replaceable Parts . . . . . 418
A5277A/AZ Controller Enclosure Specifications . . . . . 420
Dimensions: . . . . . 420
Weight: . . . . . 420
AC Power: . . . . . 421
Heat Output: . . . . . 421
Environmental Specifications . . . . . 422
Acoustics . . . . . 423
A5294A/AZ Disk Enclosure Specifications . . . . . 424
Dimensions: . . . . . 424
Weight: . . . . . 424

AC Power:. . . . . 425
DC Power Output:. . . . . 425
Heat Output:. . . . . 425
Environmental Specifications . . . . . 426
Acoustics . . . . . 427
Warranty and License Information . . . . . 428
Hewlett-Packard Hardware Limited Warranty . . . . . 428
Software Product Limited Warranty . . . . . 429
Limitation of Warranty . . . . . 429
Hewlett-Packard Software License Terms. . . . . 431
Regulatory Compliance . . . . . 434
Safety Certifications:. . . . . 434
EMC Compliance. . . . . 434
FCC Statements (USA Only) . . . . . 435
IEC Statement (Worldwide). . . . . 435
CSA Statement (For Canada Only) . . . . . 435
VCCI Statement (Japan) . . . . . 436
Harmonics Conformance (Japan) . . . . . 436
Class A Warning Statement (Taiwan) . . . . . 436
Spécification ATI Classe A (France Seulement) . . . . . 437
Product Noise Declaration (For Germany Only) . . . . . 437
Geräuschemission (For Germany Only) . . . . . 438
Declaration of Conformity . . . . . 439

# 1 PRODUCT DESCRIPTION

Product Description . . . . . 18

Disk Enclosure Components . . . . . 23

Array Controller Enclosure Components . . . . . 34

Disk Array High Availability Features. . . . . 47

Capacity Management Features . . . . . 65

# Product Description

The HP SureStore E Disk Array FC60 (Disk Array FC60) is a disk storage system that features high data availability, high performance, and storage scalability. To provide high availability, the Disk Array FC60 uses redundant, hot swappable modules, which can be replaced without disrupting disk array operation should they fail.

The Disk Array FC60 consists of two primary components: an FC60 controller enclosure, and from one to six HP SureStore E Disk System SC10 enclosures (referred to throughout this document as simply disk enclosures). The controller enclosure is responsible for providing overall control of the Disk Array FC60 by managing the communication between the host and the disk enclosures. Host communication is done through dual Fibre Channel arbitrated loops (when dual controller modules are installed). By using Fibre Channel, the Disk Array FC60 achieves a high data rate throughput. High data throughput is maintained to the disks by using up to six Ultra2 SCSI channels to the disk enclosures (one channel for each disk enclosure).

In addition to increased performance, the use of multiple disk enclosures provides scalability — simplifying the process of adding storage capacity as needed. Up to six disk enclosures can be added incrementally as storage demands increase. Each disk enclosure holds up to ten disk modules in capacities of 9.1 Gbyte, 8.2 Gbyte, 36.4 Gbyte, or 73.4 Gbyte. A fully loaded system comprising six disk enclosures, each populated with ten 73.4-Gbyte disk modules, achieves a capacity of over 3 Tbytes.

The Disk Array FC60 enclosures are designed for installation in HP's original 19-inch cabinets, which include the C2785A (1.1m), C2786A (1.6m), C2787A (2m), A1896A (1.1 m), and A1897A (1.6 m) and the HP Rack System/E Racks which includes the A490xA and A150xA Rack System/E cabinets. The Disk Array FC60 is also supported in the Rittal 9000 Series racks.

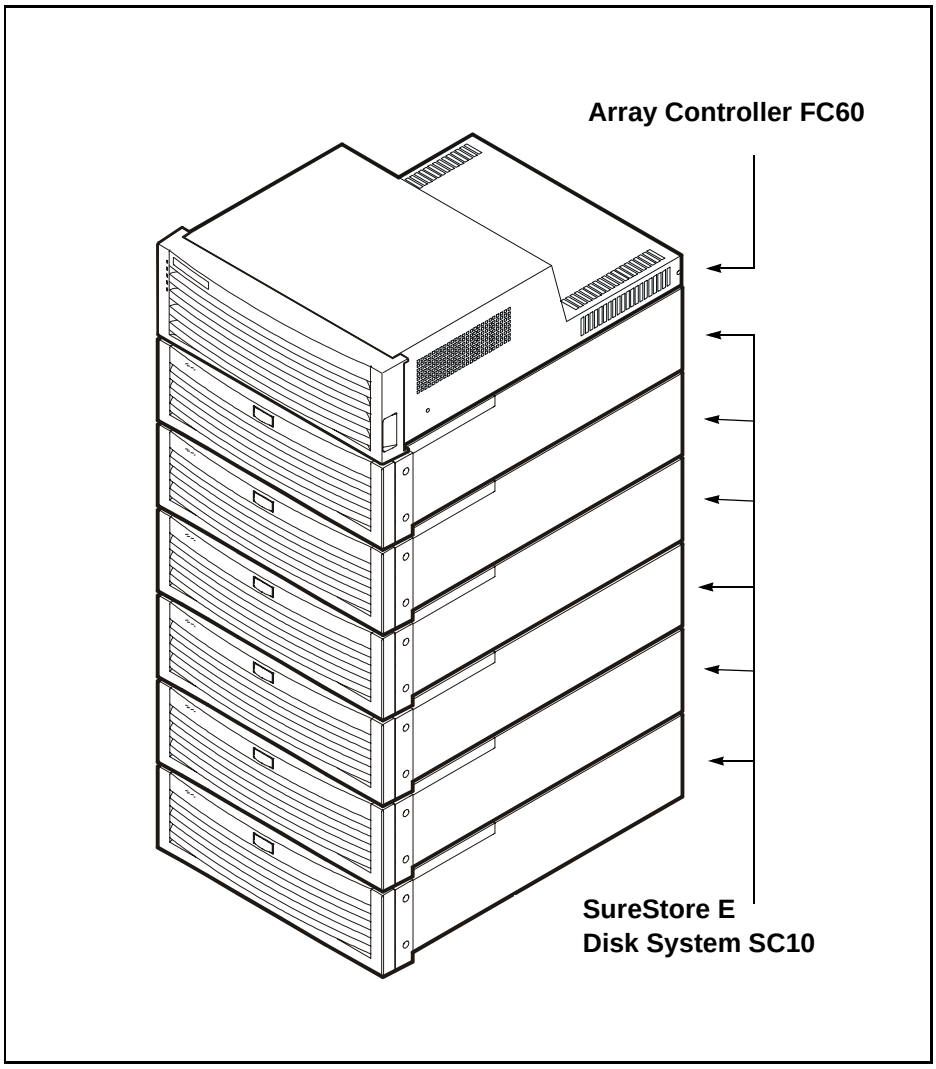



**Figure 1** HP SureStore E Disk Array FC60 (Controller with Six Disk Enclosures)

## Operating System Support

The Disk Array FC60 is currently supported on the following operating systems:

- HP-UX 11.0, 11.11, and 10.20
- Windows® NT 4.0
- Windows® 2000

**Note** Some disk array features are specific to each operating system. These features are clearly identified throughout this book.

## Management Tools

### HP-UX Tools

The following tools are available for managing the disk array on HP-UX. These tools are included with the disk array.

- Array Manager 60 command line utilities
- SAM
- Support Tools Manager (STM)

### Windows® NT and Windows® 2000

The following tools are used to manage the Disk Array FC60 on Windows NT and Windows 2000. This tool is not included with the disk array, but must be ordered separately as product A5628A.

- *HP Storage Manager 60* (A5628A)

## Features

The Disk Array FC60 offers the following features

- High availability
- Scalable storage capacity
- LED status monitoring

- RAID levels 0, 1, 0/1, 3, and 5 (RAID level 3 supported on Windows NT and Windows 2000 only)
- EMS hardware monitoring (HP-UX only)

## High Availability

High availability is a general term that describes hardware and software systems that are designed to minimize system downtime — planned or unplanned. The Disk Array FC60 qualifies as high-availability hardware, achieving 99.99% availability.

The following features enable high availability:

- Hot-swappable, high-capacity, high-speed disk modules
- Dual Fibre Channel arbitrated loops (FC-AL) connection to the host
- Redundant, hot-swappable, fans and power supplies
- Support for RAID 1, 0/1, and 5
- Remote monitoring and diagnostics
- EMS Hardware event monitoring and real-time error reporting (HP-UX only)

**Note** The Disk Array FC60 is designed to operate with either one or two controller modules; however, for data integrity and high availability it is highly recommended that dual controller modules be installed.

## Scalable Storage Capacity

The Disk Array FC60 is designed to provide maximum scalability, simplifying the process of adding storage capacity as required. Storage capacity can be added in three ways:

- By adding additional disk modules to a disk enclosure
- By adding additional disk enclosures to the array
- By replacing existing disk modules with higher capacity modules

The controller enclosure supports up to six disk enclosures. Each disk enclosure holds up to ten disk modules in capacities of 9.1 Gbyte, 18.2 Gbyte, 36.4 Gbyte, or 73.4 Gbyte. The minimum configuration for the array is one disk enclosure with two 9.1-Gbyte disk modules. The maximum configuration is six disk enclosures with ten 73.4-Gbyte disk

modules. This provides a storage capacity range from 36 Gbytes to over 3 Tbytes of usable storage.

## LED Status Monitoring

Both the controller enclosure and disk enclosure monitor the status of their internal components and operations. At least one LED is provided for each swappable module. If an error is detected in any module, the error is displayed on the appropriate module's LED. This allows failed modules to be quickly identified and replaced.

## EMS Hardware Event Monitoring (HP-UX Only)

The Disk Array FC60 is fully supported by Hewlett-Packard's EMS Hardware Monitors, which allow you to monitor all aspects of product operation and be alerted immediately if any failure or other unusual event occurs. Hardware monitoring is available at no cost to users running HP-UX 11.0 or 10.20.

Hardware monitoring provides a high level of protection against system hardware failure. It provides an important tool for implementing a high-availability strategy for your system. By using EMS Hardware Monitors, you can virtually eliminate undetected hardware failures that could interrupt system operation or cause data loss.

The EMS Hardware Monitor software with the Disk Array FC60 monitor is distributed on the *HP-UX Support Plus CD-ROM* release 9912 and later. Complete information on installing and using hardware event monitoring is contained in the EMS Hardware Monitors User's Guide (B6191-90011). A copy of this book can be accessed from the Systems Hardware, Diagnostics, & Monitoring page of Hewlett-Packard's on-line documentation web site at http://www.docs.hp.com/hpux/systems/

The minimum supported version of the Disk Array FC60 hardware monitor (fc60mon) is A.01.04. To verify the version of the monitor installed, type:

**what /usr/sbin/stm/uut/bin/tools/monitor/fc60mon**

# Disk Enclosure Components

The SureStore E Disk System SC10, or disk enclosure, is a high availability Ultra2 SCSI storage product. It provides an LVD SCSI connection to the controller enclosure and ten slots on a single-ended backplane for high-speed, high-capacity LVD SCSI disks. Six disk enclosures fully populated with 9.1 Gbtye disks provide 0.54 Tbytes of storage in a 2-meter System/E rack. When fully populated with 73.4 Gbyte disks, the array provides over 3 Tbytes of storage. These values represent maximum storage; usable storage space will vary depending on the RAID level used.

The disk enclosure consists of modular, redundant components that are easy to upgrade and maintain. See Figure 2. Disks, fans, power supplies, and Bus Control Cards (BCCs) are replaceable parts that plug into individual slots in the front and back of the disk enclosure. Redundant fans, power supplies, and disk modules can be removed and replaced without interrupting storage operations. In addition, a single disk within a LUN can be replaced while the system is on.

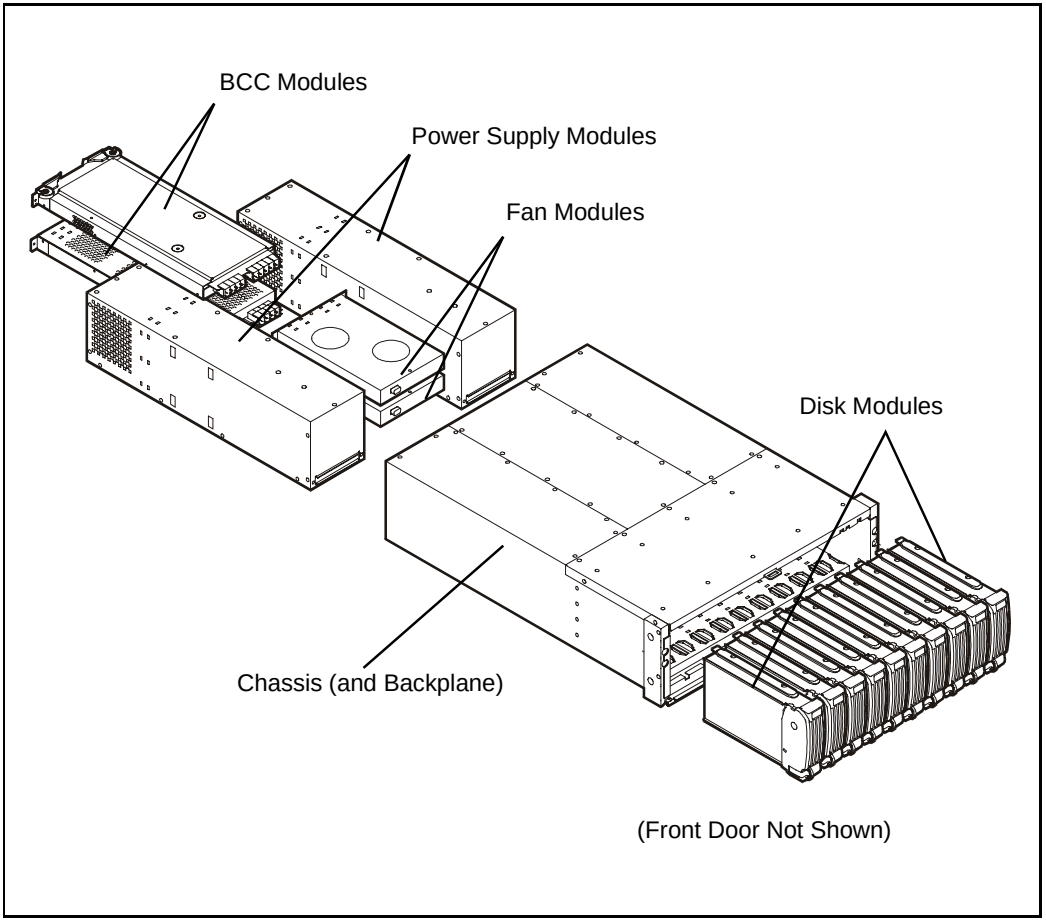



**Figure 2** Disk Enclosure Components, Exploded View

## Operation Features

The disk enclosure is designed to be installed in a standard 19-inch rack and occupies 3.5 EIA units (high). Disk drives mount in the front of the enclosure. Also located in the front of the enclosure are a power switch and status LEDs. A lockable front door shields RFI and restricts access to the disk drives and power button (Figure 3 on page 26).

BCCs are installed in the back of the enclosure along with redundant power supplies and fans.

### Status Indicators

LEDs on the front and back of the disk enclosure enable you to quickly identify and replace failed components, thereby preventing or minimizing downtime. "Troubleshooting" on page 359 provides more detailed information about the operation of these LEDs.

Two system LEDs on the front, top right corner of the disk enclosure (A in Figure 3) indicate the status of the disk enclosure. The left LED indicates when power is on or off and the right LED identifies if a fault has occurred. Additional pairs of LEDs above each disk slot (D in Figure 3), indicate disk activity and a fault condition. The left LED (green) indicates disk I/O activity and the right LED goes on if the disk module has experienced a fault. The disk fault LEDs are also used by the management tools to identify a specific disk module by flashing its fault LED.

On the back of the disk enclosure, the following LEDs (K) indicate the status of replaceable components and Fibre Channel link:

- Power supply status and fault LEDs
- Fan status and fault LEDs

Bus Controller Card LEDs:

- BCC Fault LED
- Term Power LED (monitors power on the SCSI bus)
- Full Bus Mode LED
- LVD Mode LED
- Bus Free Status LED

For detailed information on LED operation, refer to "Troubleshooting" on page 359

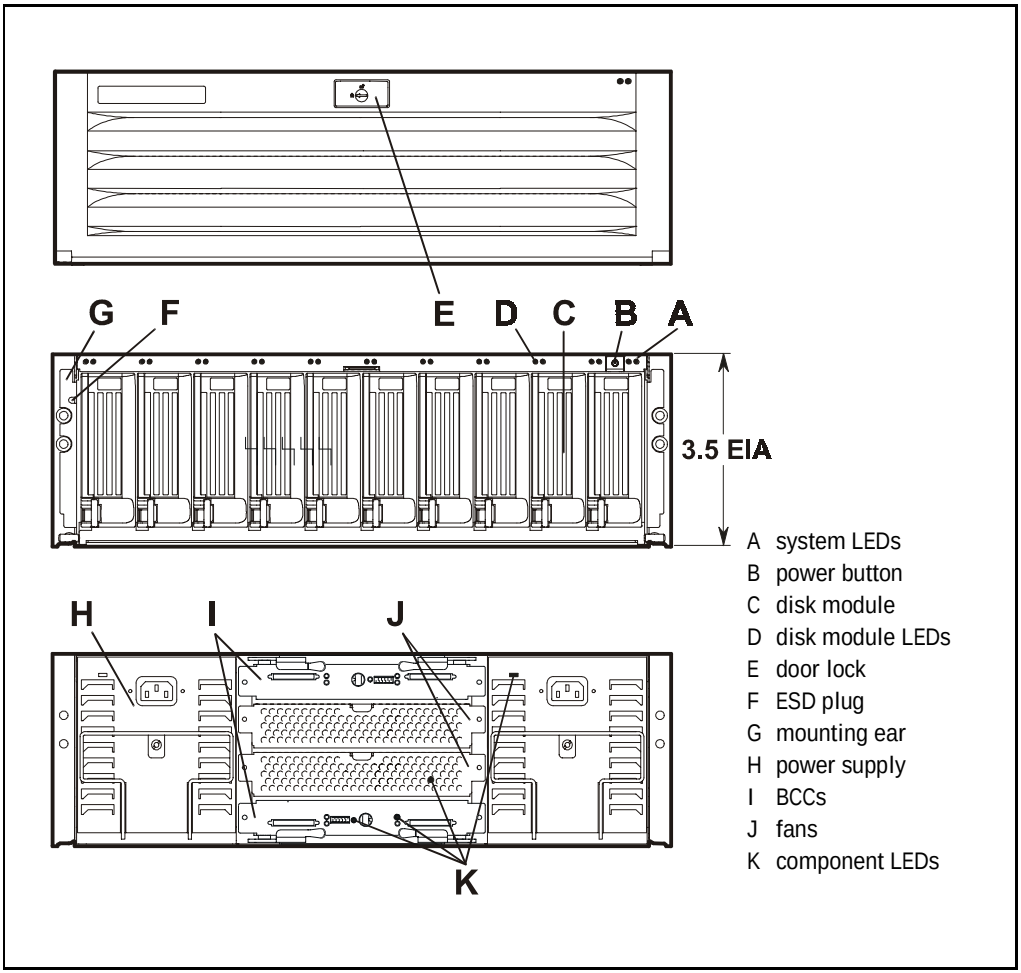



**Figure 3** Disk Enclosure Front and Back View

## Power Switch

The power switch (B in Figure 3) interrupts power from the power supplies to the disk enclosure components. Power to the power supplies is controlled by the power cords and the AC source.

## Disk Enclosure SC10 Modules

The disk enclosure hot-swappable modules include the following:

- Disks and fillers
- Fans
- Power supplies

### Disks and Fillers

Hot-swappable disk modules make it easy to add or replace disks. Fillers are required in all unused slots to maintain proper airflow within the enclosure.

Figure 4 illustrates the 3.5-inch disks in a metal carrier. The open carrier design allows ten half height (1.6 inch) disks to fit the 19-inch width of a standard rack and meet cooling needs.

**WARNING** **Touching exposed circuits on the disk module can damage the disk drive inside. To avoid damage, always handle disks carefully and use ESD precautions.**

The following plastic parts of the disk are safe to touch:

- Bezel-handle (A in Figure 4)
- Cam latch (B)
- Insertion guide (F)

Metal standoffs (D) protect exposed circuits against damage when the disk is laid circuit-side down on a flat surface.

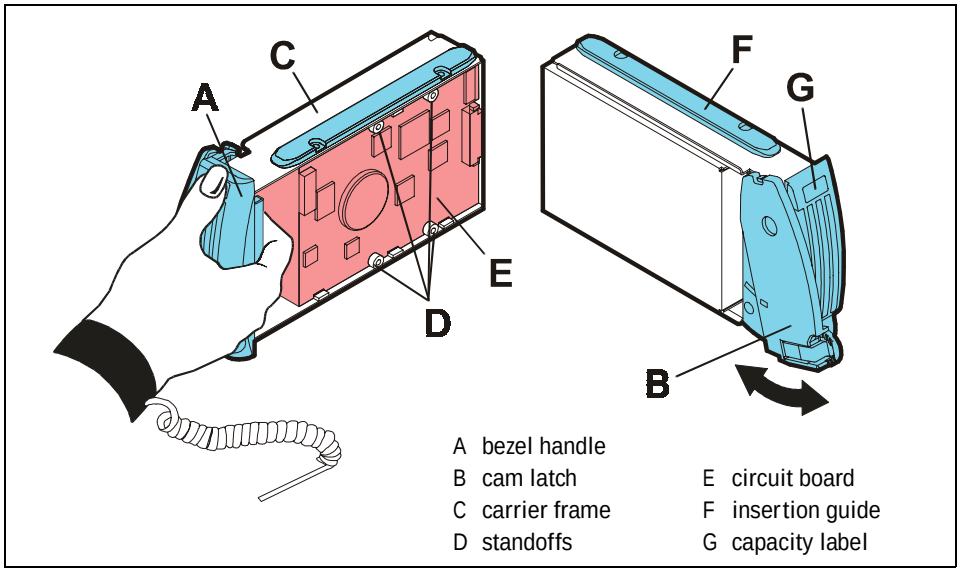



**Figure 4** Disk Module

Disks fit snugly in their slots. The cam latch (B in Figure 4) is used to seat and unseat the connectors on the backplane.

A label (G) on the disk provides the following information:

- Disk mechanism height: 1.6 inch (half height) or 1 inch (low profile)
- Rotational speed: 10K RPM and 15K RPM (18 Gbyte only)
- Capacity: 9.1 Gbyte, 18.2 Gbyte, 36.4 Gbyte, or 73.4 Gbyte

A large zero on the capacity label distinguishes a filler from a disk. Fillers are required in all unused slots to maintain proper airflow within the enclosure.

***CAUTION*** *Fillers must be installed in unused slots to maintain proper cooling within the disk enclosure.*

## BCCs

Two Backplane Controller Cards, BCCs, control the disks on one or two buses according to the setting of the Full Bus switch. When the Full Bus switch is set to on, BCC A, in the top slot, accesses the disks in all ten slots. When the Full Bus switch is off, BCC A accesses disks in the even-numbered slots and BCC B accesses disks in the odd-numbered slots.

**Note** In full bus mode, all ten disks can be accessed through either BCC. However, internally each BCC still manages five disks. This means that if the BCC that is not connected to the SCSI channel fails, access to its five disks will be lost. Failure of the BCC that is connected to the SCSI channel will render all ten disks inaccessible.

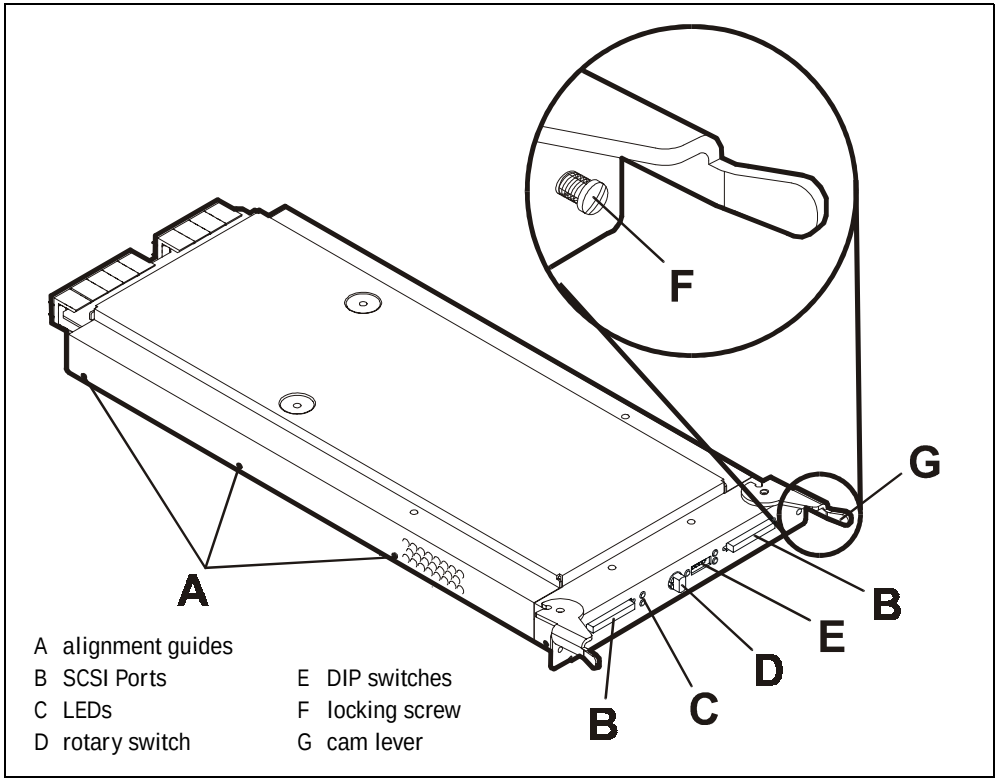



**Figure 5** BCC

Each BCC provides two LVD SCSI ports (B in Figure 5) for connection to the controller enclosure.

The EEPROM on each BCC stores 2 Kbytes of configuration information and user-defined data, including the manufacturer serial number, World Wide Name, and product number.

The following are additional features of the BCC:

- LEDs (C in Figure 5) show the status of the BCC and the bus.
- A rotary switch (D) used to set the enclosure (tray) ID which is used by internal controller operations, and also by the management tools to identify each enclosure.
- DIP switches set disk enclosure options. The only option used by the Disk Array FC60 is the full-bus/split-bus mode.
- Screws (F) prevent the card from being unintentionally disconnected.
- Cam levers (G) assist in installing and removing the BCC from the enclosure, ensuring a tight connection with the backplane.

BCC functions include drive addressing, fault detection, and environmental services.

## Fans

Redundant, hot-swappable fans provide cooling for all enclosure components. Each fan has two internal high-speed blowers (A in Figure 6), an LED (B), a pull tab (C), and two locking screws (D).

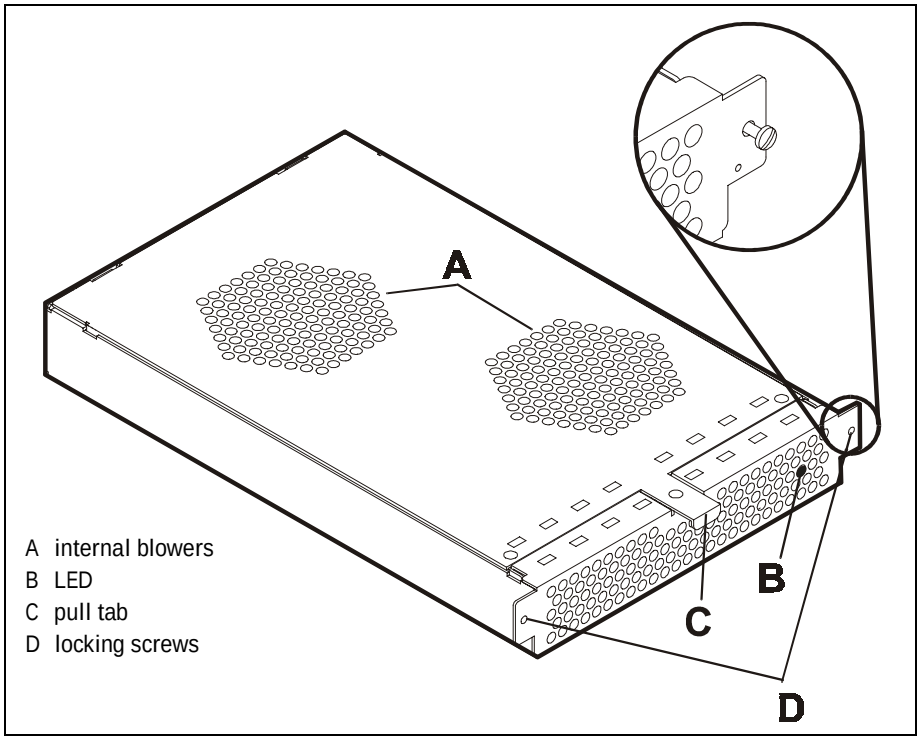



**Figure 6** Fan

Internal circuitry senses blower motion and triggers a fault when the speed of either blower falls below a critical level. If a fan failure occurs, the amber fault LED will go on. An alert should also be generated by EMS Hardware Monitoring when a fan failure occurs.

## Power Supplies

Redundant, hot-swappable 450-watt power supplies convert wide-ranging AC voltage from an external main to stable DC output and deliver it to the backplane. Each power supply has two internal blowers, an AC receptacle (A in Figure 7), a cam handle (B) with locking screw, and an LED (C). Internal control prevents the rear DC connector from becoming energized when the power supply is removed from the disk enclosure.

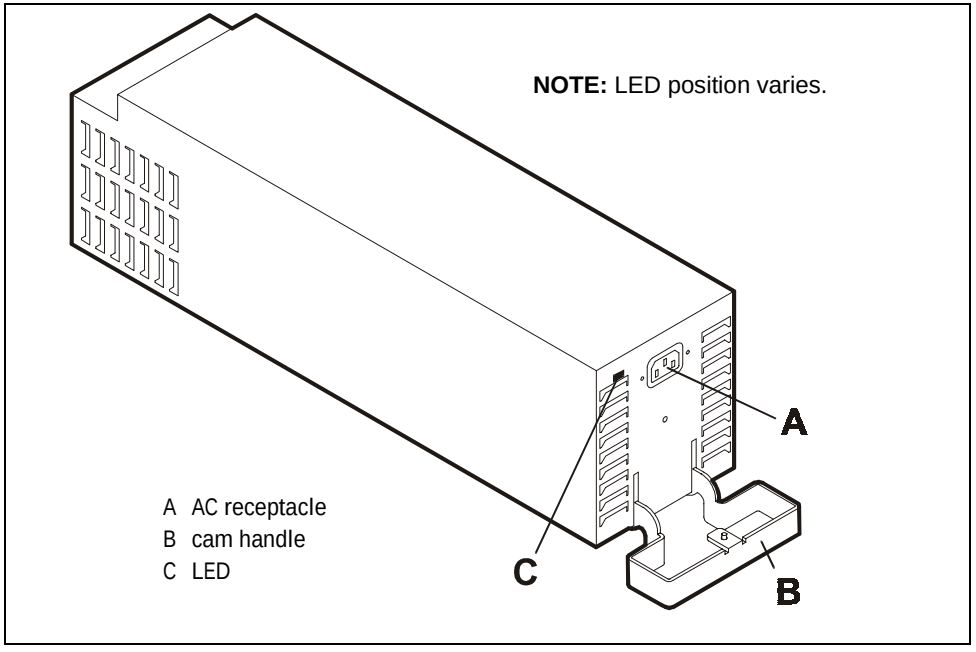



**Figure 7** Power Supply

**Note** Although it is possible to operate the disk enclosure on one power supply, it is not recommended. Using only one supply creates a single point of failure. If the power supply fails, the entire enclosure is inaccessible. To maintain high availability, both power supplies should be used at all times, and a failed supply should be replaced as soon as possible.

Power supplies share the load reciprocally; that is, each supply automatically increases its output to compensate for reduced output from the other. If one power supply fails, the other delivers the entire load.

Internal circuitry triggers a fault when a power supply fan or other power supply part fails. If a power supply failure occurs, the amber fault LED will go on. An alert should also be generated by EMS Hardware Monitoring when a power supply failure occurs.

# Array Controller Enclosure Components

The array controller enclosure, like the disk enclosure, consists of several modules that can be easily replaced, plus several additional internal assemblies. See Figure 8. Together, these removable modules and internal assemblies make up the field replaceable units (FRUs). Many modules can be removed and replaced without disrupting disk array operation.

The following modules are contained in the controller enclosure:

- Controller modules
- Controller fan module
- Power supply modules
- Power supply fan module
- Battery backup unit

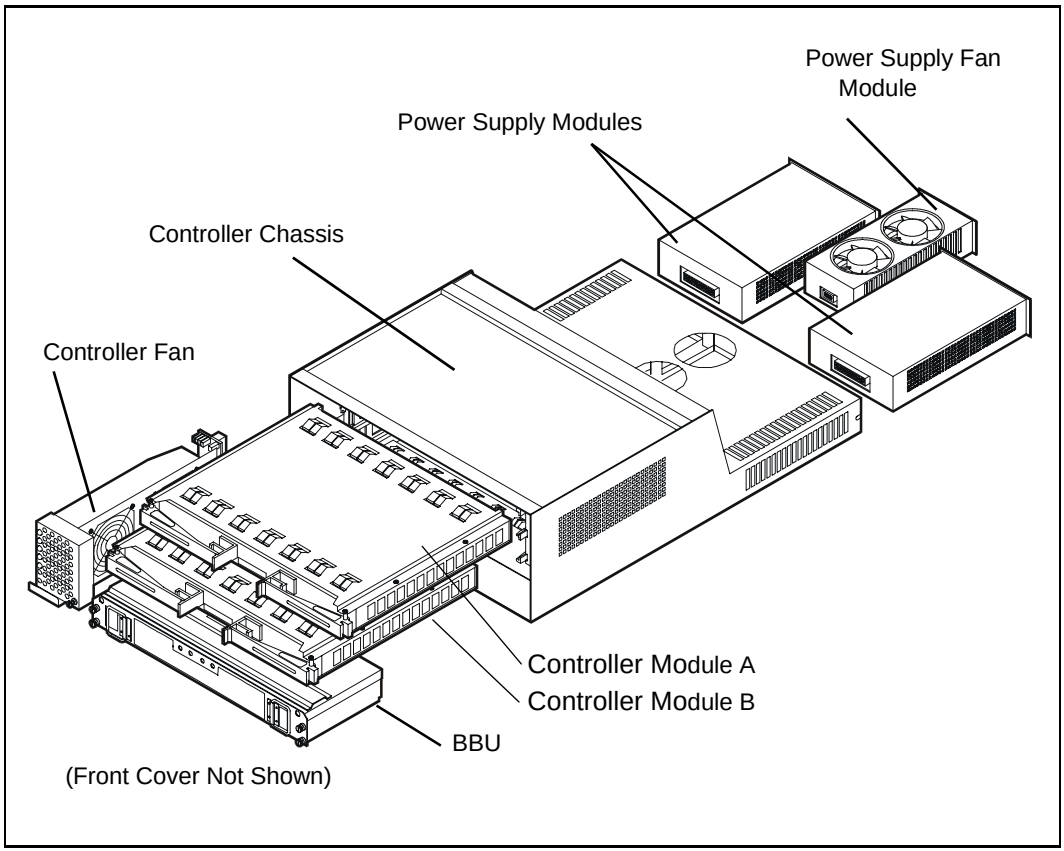



**Figure 8** Controller Enclosure Exploded View

During operation, controller enclosure status is indicated by five LEDs on the front left of the controller enclosure. Faults detected by the controller module cause the corresponding controller enclosure fault LED to go on. Additional LEDs on the individual components identify the failed component. See "Troubleshooting" on page 359 for detailed information on LED operation.

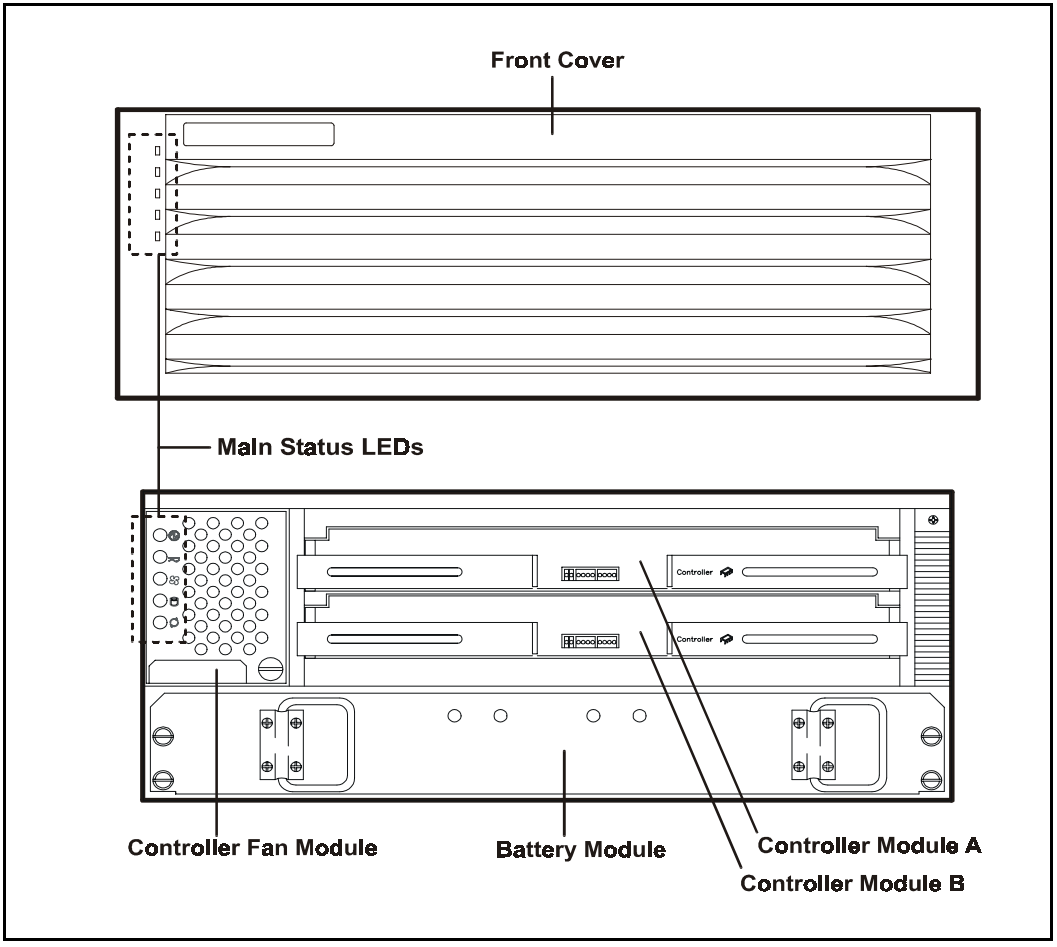



**Figure 9** Controller Enclosure Front View

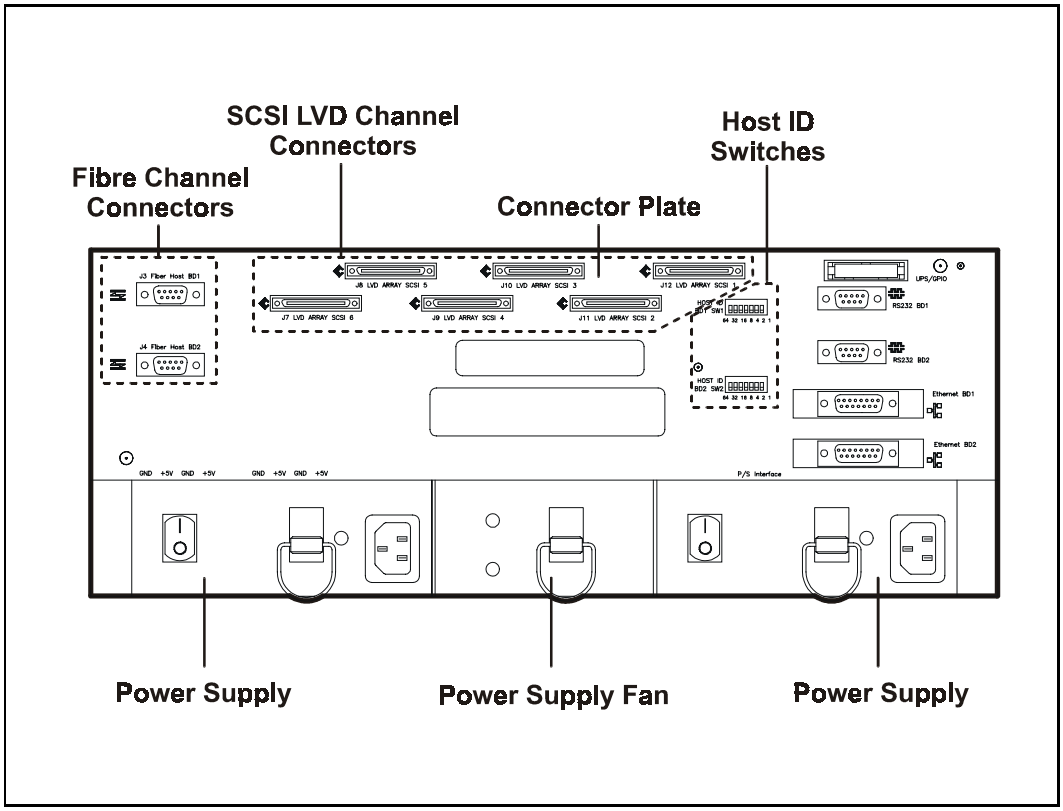



**Figure 10** Controller Enclosure Rear View

## Front Cover

The controller enclosure has a removable front cover which contains slots for viewing the main operating LEDs. The cover also contains grills that aid air circulation. The controller modules, controller fan, and battery backup unit are located behind this cover. This cover must be removed to gain access to these modules, and also, to observe the controller status and BBU LEDs.

## Controller Modules

The controller enclosure contains one or two controller modules. See Figure 11. These modules provide the main data and status processing for the Disk Array FC60. The controller modules slide into two controller slots (A and B) and plug directly into the backplane. Two handles lock the modules in place. Each controller slot has a controller letter that identifies the physical location of the controller in the chassis: controller slot A or controller slot B (also known as BD1 and BD2, respectively, as referenced on the back of the controller enclosure).

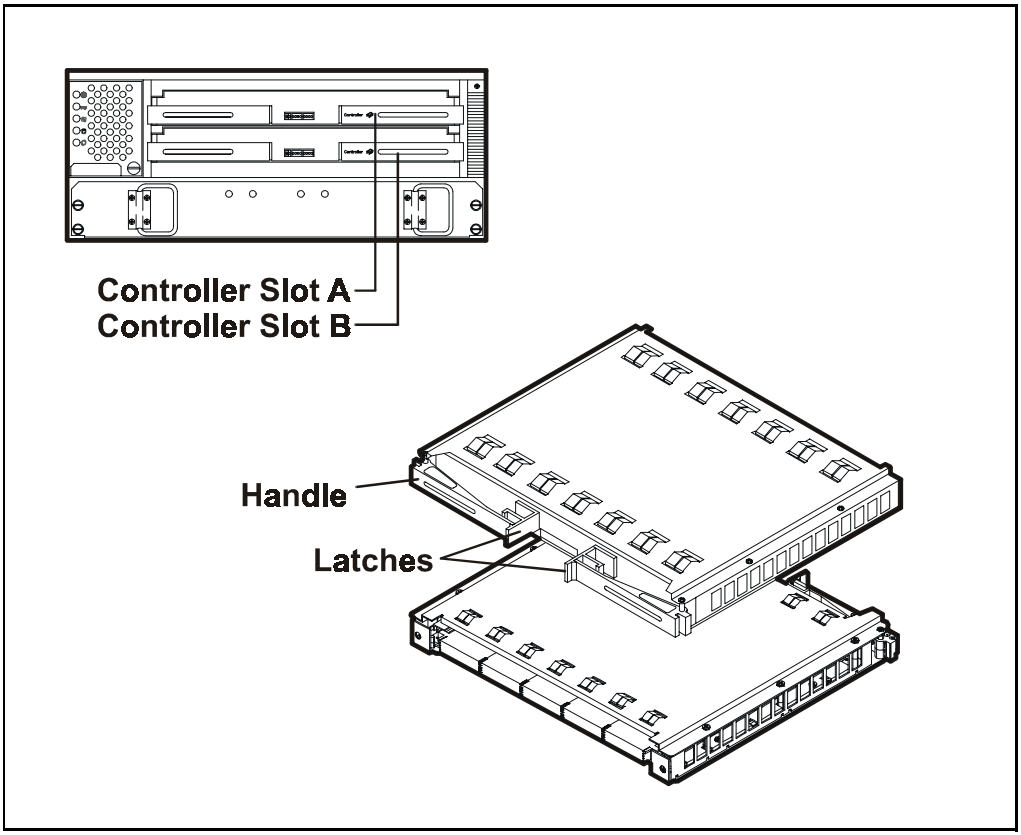



**Figure 11** Controller Modules

Each controller module has ten LEDs. See Figure 12. One LED identifies the controller module's power status. A second LED indicates when a fault is detected. The remaining eight LEDs provide detailed fault condition status. The most significant LED, the heartbeat, flashes approximately every two seconds beginning 15 seconds after power-on. "Troubleshooting" on page 359 contains additional information on controller LED operation.

The controller module connects to the host via Fibre Channel, and to the disk enclosures via LVD SCSI. Each controller must have a unique host fibre ID number assigned using the ID switches on the back of the controller modules. See "Installation" on page 143 for more information on setting host IDs.

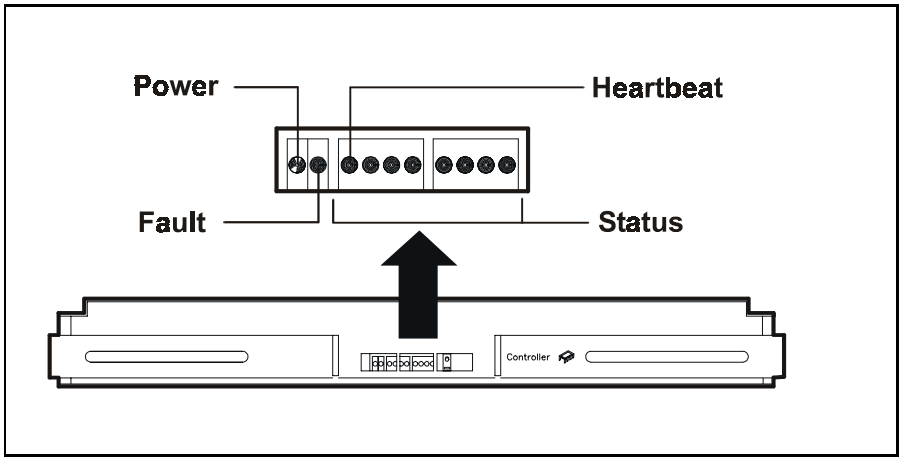



**Figure 12** Controller Module LEDs

### Controller Memory Modules

Each controller module contains SIMM and DIMM memory modules. Two 16-Mbyte SIMMs (32 Mbytes total) store controller program and other data required for operation. The standard controller module includes 256-Mbytes of cache DIMM, which is upgradeable to 512 Mbytes. The cache may be configured as either two 128-Mbyte DIMMs, or a single 256-Mbyte DIMM. Cache memory serves as temporary data storage during read and write operations, improving I/O performance. When cache memory contains write data, the Fast Write Cache LED, on the front of the controller enclosure is on. See Figure 13.

## Controller Fan Modules

The controller fan module is a single removable unit containing dual cooling fans and temperature monitoring logic. See Figure 13. It includes five LEDs that indicate overall system status and controller fan status. The fans provide cooling by pulling air in through ventilation holes, moving it across the controller cards, and exhausting it out the ventilation holes in the fan assembly. The dual fans provide a redundant cooling system to both controller modules. If one fan fails, the other continues to operate and provide sufficient air circulation to prevent the controllers from overheating until the fan module is replaced. The fan module plugs into a slot on the front of the controller enclosure, and has a handle and captive screw for easy service.

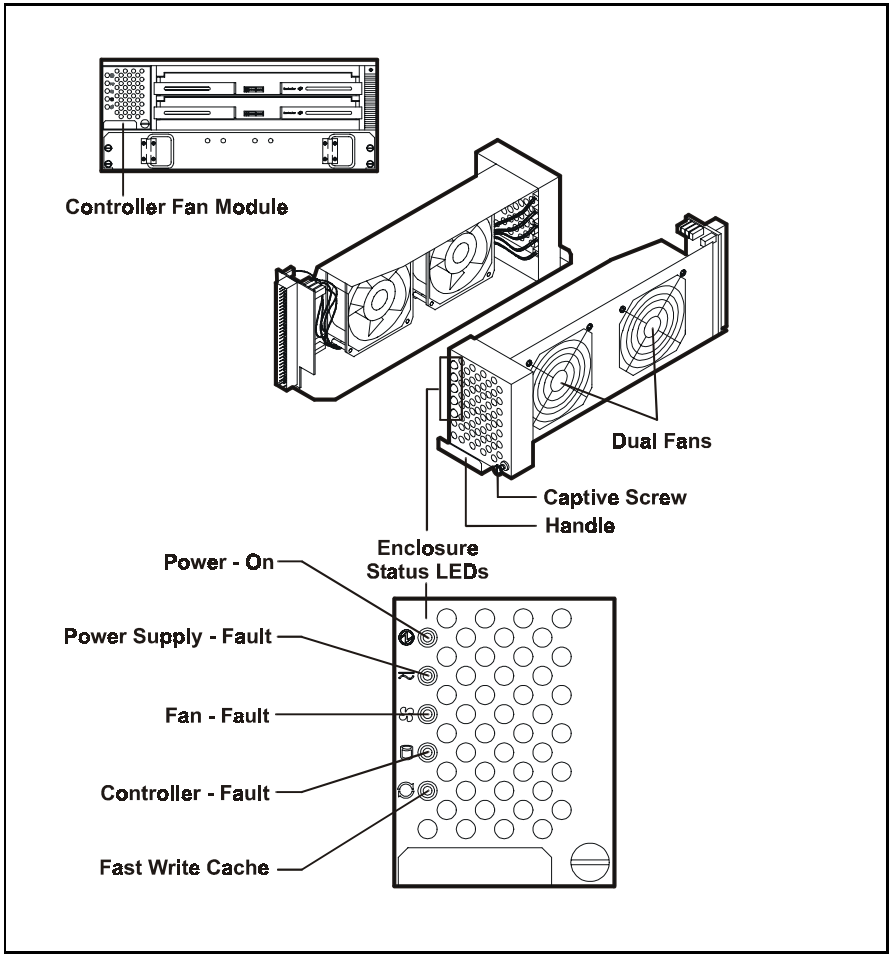



**Figure 13** Controller Fan Module

## Power Supply Modules

Two separate power supplies provide electrical power to the internal components by converting incoming AC voltage to DC voltage. Both power supplies are housed in removable power supply modules that slide into two slots in the back of the controller and plug directly into the power interface board. See Figure 14.

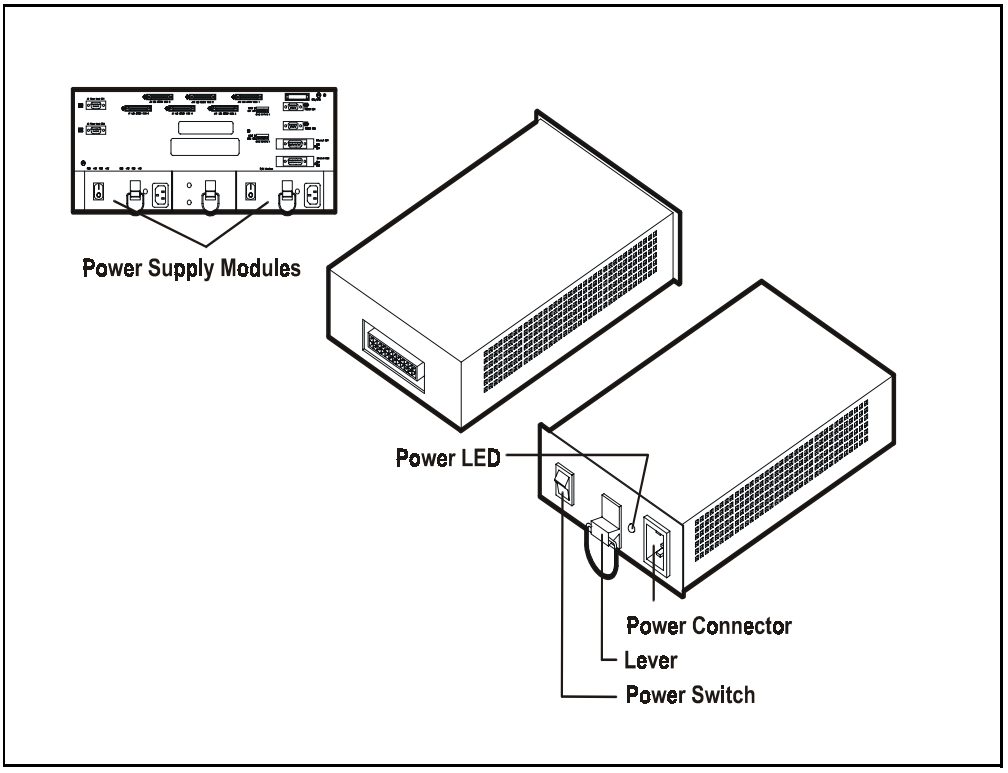



**Figure 14** Power Supply Modules

Each power supply uses a separate power cord. These two power cords are special ferrite bead cords (part no. 5064-2482) required for FCC compliance. Both power cords can plug into a common power source or each cord can plug into a separate circuit (to provide power source redundancy).

Each power supply is equipped with a power switch to disconnect power to the supply. Turning off both switches turns off power to the controller. This should not be performed unless I/O activity to the disk array has been stopped, and the write cache has been flushed as indicated by the Fast Write Cache LED being off.

***CAUTION*** *The controller power switches should not be turned off unless all I/O activity to the disk array has been suspended from the host. Also, power should not be turned off if the Fast Write Cache LED is on, indicating that there is data in the write cache. Wait until the Fast Write Cache LED goes off before shutting off power to the disk array.*

Each power supply is equipped with a power-on LED indicator. If the LED is on (green) the supply is providing dc power to the controller. If the LED is off, there is a malfunction or the power has been interrupted. The system Power Fault LED on the front of the controller enclosure works in conjunction with the Power Supply LEDs. If both power supplies are on, the system Power Fault LED will be off. If either power supply is off or in a fault state, the system Power Fault LED goes on. When both power supplies are off or not providing power to the enclosure, the system power LED on the front of the controller enclosure will be off.

## Power Supply Fan Module

Like the controller fan, the power supply fan module (Figure 15) is a single removable unit that contains dual cooling fans. Dual fans provide a redundant cooling system to both power supply modules. If one fan fails, the other will continue to operate. A single fan will provide sufficient air circulation to prevent the power supplies from overheating until the entire power supply fan module can be replaced. The power supply fan module plugs directly into a slot on the back of the controller enclosure, between the power supplies. It has a locking lever that allows it to be unlatched and removed.

The power supply fan can be hot swapped, provided the exchange is performed within 15 minutes. This time limit applies only to the total time during which the fan is out of the enclosure, beginning when you remove the failed unit and ending when you re-seat the new one. This does not include the time it takes you to perform this entire procedure (including checking LEDs).

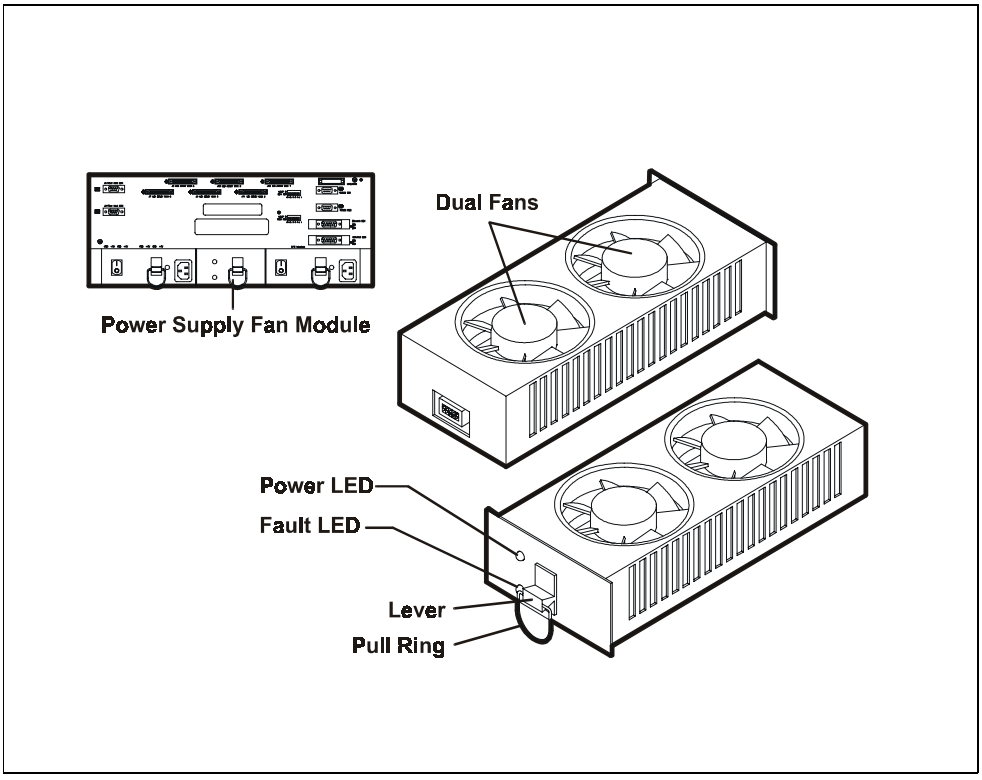



**Figure 15** Power Supply Fan Module

## Battery Backup Unit

The controller enclosure contains one removable battery backup unit (BBU) that houses two rechargeable internal batteries (A and B) and a battery charger board. The BBU plugs into the front of the controller enclosure where it provides backup power to the controller's cache memory during a power outage. The BBU will supply power to the controllers for up to five days (120 hrs). All data stored in memory will be preserved as long as the BBU supplies power. When power to the disk array is restored, the cache data will be written to disk.

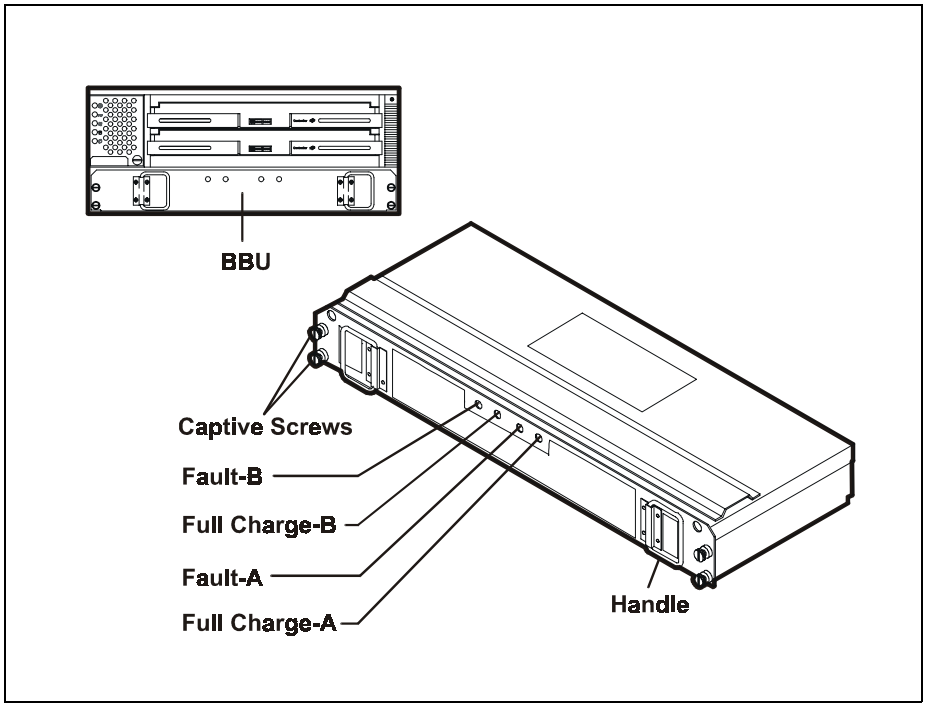



**Figure 16** Battery Backup Unit

***CAUTION*** *During a power outage, do not remove the controller or the BBU. Removing either of these modules can compromise data integrity.*

The BBU contains four LEDs that identify the condition of the battery. Internally, the BBU consists of two batteries or banks, identified as bank “A” and bank “B.” During normal operation both of the Full Charge LEDs (Full Charge-A and Full Charge-B) are on and the two amber Fault LEDs are off. If one or both of the Fault LEDs are on, refer to "Troubleshooting" on page 359 for information on solving the problem. The Full Charge LEDs flash while the BBU is charging. It can take up to seven hours to fully charge a new BBU.

### Battery Operation and Replacement

Replace the BBU every two years or whenever it fails to hold a charge, as indicated by the BBU Fault LEDs. The service label on the BBU provides a line for recording the date on which the BBU was serviced. Check this label to determine when to replace the BBU. When a BBU is replaced, it may require up to seven hours to fully charge. The Full Charge LEDs flash while the BBU is charging, and remain on when the BBU is fully charged.

If you replace the BBU and still experience battery-related problems (such as a loss of battery power to the controllers or batteries not charging properly), the controller enclosure may have some other internal component failure. In this case contact your HP service engineer.

### Battery Operation for No Data Loss

The BBU protects the write cache (data which has not been written to disk) for at least 120 hours (five days) in case of a power failure. When power to the disk array is restored, data in the cache is written to the disks and no data loss occurs. However, if the system is to be powered off for more than 120 hours, it is imperative that a proper shutdown procedure be executed or data may be lost. The following are recommendations:

- Battery status must always be checked and replaced when a failure is indicated.
- Never remove the BBU without first performing a proper shutdown procedure.
- For a planned shutdown, make sure that all data has been written to disks before removing power. This is indicated by the Fast Write Cache LED which will be off when there is no longer any data in write cache. See Figure 13.
- If the BBU is removed, do not shut off power to the array unless the Fast Write Cache LED is off. Data in write cache will be posted to disk 10 seconds after the BBU is removed.

# Disk Array High Availability Features

High availability systems are designed to provide uninterrupted operation should a hardware failure occur. Disk arrays contribute to high availability by ensuring that user data remains accessible even when a disk or other component within the Disk Array FC60 fails. Selecting the proper Fibre Channel topology and system configuration can protect against the failure of any hardware component in the I/O path to the disk array by providing an alternate path to all user data.

The Disk Array FC60 provides high availability in the following ways:

- Supported RAID levels 1, 0/1, 3, and 5 all use data redundancy to protect data when a disk failure occurs. RAID 0 is supported but it does not offer data redundancy and should not be used in high-availability environments.
- Global hot spare disks serve as automatic replacements for failed disks.
- Alternate hardware paths to user data protects against I/O path failures.
- Redundant, hot-swappable hardware components can be replaced without interrupting disk array operation.

## RAID Technology

RAID technology contributes to high availability through the use of data redundancy, which ensures that data on the disk array remains available even if a disk or channel failure occurs. RAID technology uses two techniques to achieve data redundancy: mirroring and parity. A third characteristic of RAID technology, data striping, enhances I/O performance.

## Disk Mirroring

Disk mirroring achieves data redundancy by maintaining a duplicate copy of all data. Disks are organized into pairs: one disk serves as the data disk, the other as the mirror which contains an exact image of its data. If either disk in the pair fails or becomes inaccessible, the remaining disk provides uninterrupted access to the data.

The disk array uses hardware mirroring, in which the disk array automatically synchronizes the two disk images, without user or operating system involvement. This is unlike the software mirroring, in which the host operating system software (for example, LVM) synchronizes the disk images.

Disk mirroring is used by RAID 1 and RAID 0/1 LUNs. A RAID 1 LUN consists of exactly two disks: a primary disk and a mirror disk. A RAID 0/1 LUN consists of an even number of disks, half of which are primary disks and the other half are mirror disks. If a disk fails or becomes inaccessible, the remaining disk of the mirrored pair provides uninterrupted data access. After a failed disk is replaced, the disk array automatically rebuilds a copy of the data from its companion disk. To protect mirrored data from a channel or internal bus failure, each disk in the LUN should be in a different enclosure.

## Data Parity

Data parity is a second technique used to achieve data redundancy. If a disk fails or becomes inaccessible, the parity data can be combined with data on the remaining disks in the LUN to reconstruct the data on the failed disk. Data parity is used for RAID 3 and RAID 5 LUNs.

To ensure high availability, each disk in the LUN should be in a separate enclosure. Parity cannot be used to reconstruct data if more than one disk in the LUN is unavailable.

Parity is calculated on each write I/O by doing a serial binary exclusive OR (XOR) of the data segments in the stripe written to the data disks in the LUN. The exclusive OR algorithm requires an even number of binary 1s to create a result of 0.

Figure 17 illustrates the process for calculating parity on a five-disk LUN. The data written on the first disk is “XOR’d” with the data written on the second disk. The result is “XOR’d” with the data on the third disk, which is “XOR’d” with the data on the fourth disk. The result, which is the parity, is written to the fifth disk. If any bit changes state, the parity also changes to maintain a result of 0.

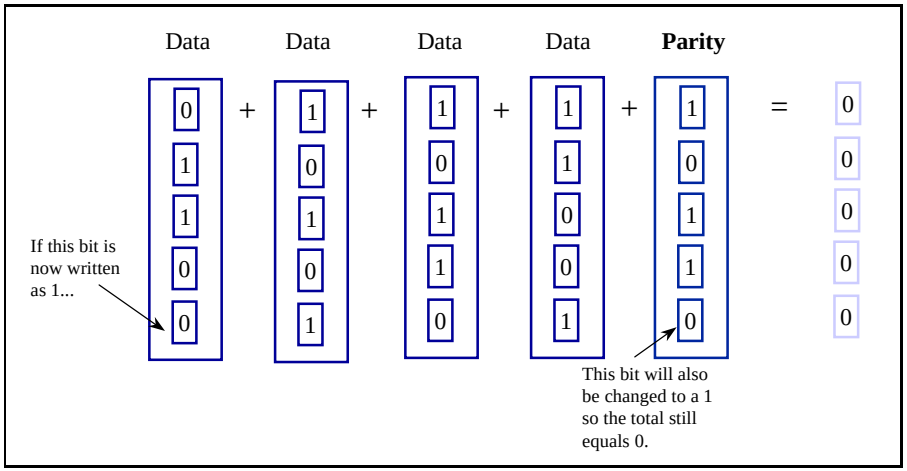



**Figure 17** Calculating Data Parity

## Data Striping

Data striping, which is used on RAID 0, 0/1, 3 and 5 LUNs, is the performance-enhancing technique of reading and writing data to uniformly sized segments on all disks in a LUN simultaneously. Collectively, the segments comprise a stripe of data on the LUN. Data striping enhances performance by allowing multiple sets of read/write heads to execute the same I/O transaction simultaneously.

The amount of information simultaneously read from or written to each disk is the stripe segment size. The stripe segment size is configurable to provide optimum performance under varying sizes of I/O transactions. Stripe segment size is specified in 512-byte blocks of data.

Stripe segment size can affect disk array performance. The smaller the stripe segment size, the more efficient the distribution of data read or written across the stripes in the LUN. However, if the stripe segment is too small for a single I/O operation, the operation requires access to two stripes. Called a stripe boundary crossing, this action may negatively impact performance.

The optimum stripe segment size is the smallest size that will rarely force I/Os to a second stripe. For example, assume your application uses a typical I/O size of 64 KB. If you are

using a 5-disk RAID 5 LUN, a stripe segment size of 32 blocks (16 KB) would ensure that an entire I/O would fit on a single stripe (16 KB on each of the four data disks).

The total stripe size is the number of disks in a LUN multiplied by the stripe segment size. For example, if the stripe segment size is 32 blocks and the LUN comprises five disks, the stripe size is 32 X 5, or 160 blocks (81,920 bytes).

## RAID Levels

RAID technology uses a number of different techniques for storing data and maintaining data redundancy. These industry-standard RAID levels define the method used for distributing data on the disks in a LUN. LUNs that use different RAID levels can be created on the same disk array.

The Disk Array FC60 supports the following RAID levels:

- RAID 0
- RAID 1
- RAID 0/1
- RAID 3 (Windows NT and Windows 2000 only)
- RAID 5

### RAID 0

***CAUTION*** *RAID 0 does not provide data redundancy. It should only be used in situations where high performance is more important than data protection. The failure of any disk within a RAID 0 LUN will cause the loss of all data on the LUN.. RAID 0 should only be used for non-critical data that could be lost in the event of a hardware failure.*

RAID 0 uses disk striping to achieve high performance. Data is striped across all disk in the LUN. The ability to access all disks in the LUN simultaneously provides a high I/O rate. A RAID 0 group configuration for a logical disk unit offers fast access, but without the high availability offered by the other RAID levels.

Unlike other RAID levels, RAID 0 does not provide data redundancy, error recovery, or other high availability features. Consequently it should not be used in environments where high-availability is critical. All data on a RAID 0 LUN is lost if a single disk within the LUN

fails. RAID-0 provides enhanced performance through simultaneous I/Os to multiple disk modules. Software mirroring the RAID-0 group provides high availability.

Figure 18 illustrates the distribution of user and parity data in a four-disk RAID 0 LUN. The the stripe segment size is 8 blocks, and the stripe size is 32 blocks (8 blocks times 4 disks). The disk block addresses in the stripe proceed sequentially from the first disk to the second, third, and fourth, then back to the first, and so on.

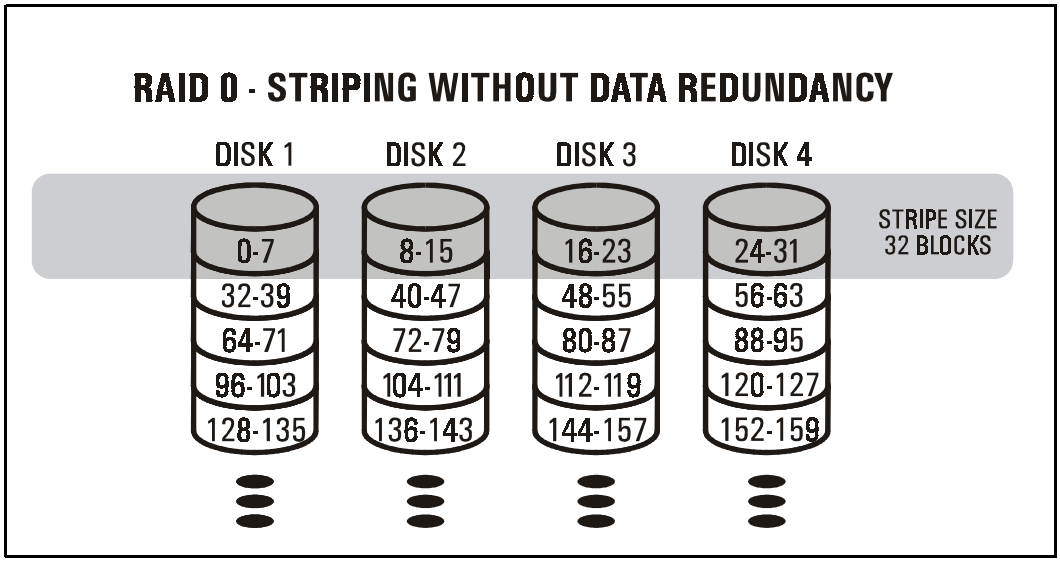



**Figure 18** RAID 0 LUN

## RAID 1

RAID 1 uses mirroring to achieve data redundancy. RAID 1 provides high availability and good performance, but at the cost of storage efficiency. Because all data is mirrored, a RAID 1 LUN has a storage efficiency of 50%.

A RAID 1 LUN consists of exactly two disks configured as a mirrored pair. One disk is the data disk and the other is the disk mirror. The disks in a RAID 1 LUN are mirrored by the disk array hardware, which automatically writes data to both the data disk and the disk mirror. Once bound into a RAID 1 mirrored pair, the two disks cannot be accessed as

individual disks. For highest data availability, each disk in the mirrored pair must be located in a different enclosure.

When a data disk or disk mirror in a RAID 1 LUN fails, the disk array automatically uses the remaining disk for data access. Until the failed disk is replaced (or a rebuild on a global hot spare is completed), the LUN operates in degraded mode. While in degraded mode the LUN is susceptible to the failure of the second disk. If both disks fail or become inaccessible simultaneously, the data on the LUN becomes inaccessible.

Figure 19 shows the distribution of data on a RAID 1 LUN. Note that all data on the data disk is replicated on the disk mirror.

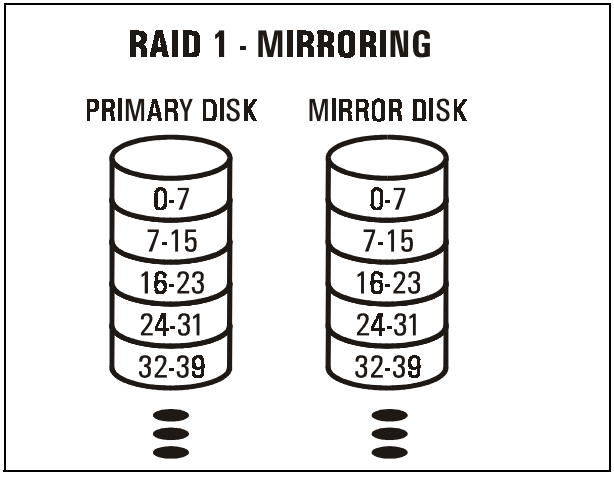



**Figure 19** RAID 1 LUN

## RAID 0/1

RAID 0/1 uses mirroring to achieve data redundancy and disk striping to enhance performance. It combines the speed advantage of block striping with the redundancy advantage of mirroring. Because all data is mirrored, a RAID 0/1 LUN has a storage efficiency of 50%.

A RAID 0/1 LUN contains an even number of from four to 30 disks. One half of the disks are primary disks and the other half are disk mirrors. The disks in a RAID 0/1 LUN are mirrored by the disk array hardware, which automatically writes data to both disks in the mirrored

pair. For highest data availability, each disk in the mirrored pair must be located in a different enclosure.

When a disk fails, the disk array automatically uses the remaining disk of the mirrored pair for data access. A RAID 0/1 LUN can survive the failure of multiple disks, as long as one disk in each mirrored pair remains accessible. Until the failed disk is replaced (or a rebuild on a global hot spare is completed), the LUN operates in degraded mode. While in degraded mode, the LUN is susceptible to the failure of the second disk of the pair. If both disks fail or become inaccessible simultaneously, the data on the LUN becomes inaccessible.

Figure 20 illustrates the distribution of data in a four-module RAID 0/1 LUN. The disk block addresses in the stripe proceed sequentially from the first pair of mirrored disks (disks 1 and 2) to the second pair of mirrored disks (disks 3 and 4), then again from the first mirrored disks, and so on.

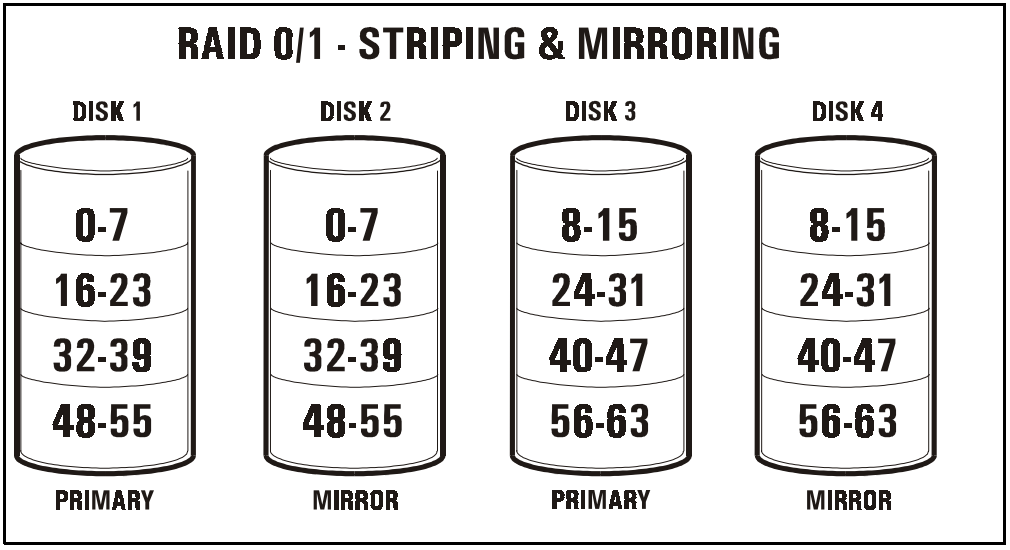



**Figure 20** RAID 0/1 LUN

## RAID 3

RAID 3 uses parity to achieve data redundancy and disk striping to enhance performance. Data is distributed across all but one of the disks in the RAID 3 LUN. The remaining disk is used to store parity information for each data stripe. A RAID 3 LUN consists of three or

more disks. For highest availability, the disks in a RAID 3 LUN must be in different enclosures.

If a disk fails or becomes inaccessible, the disk array can dynamically reconstruct all user data from the data and parity information on the remaining disks. When a failed disk is replaced, the disk array automatically rebuilds the contents of the failed disk on the new disk. The rebuilt LUN contains an exact replica of the information it would have contained had the disk not failed.

Until a failed disk is replaced (or a rebuild on a global hot spare is completed), the LUN operates in degraded mode. The LUN must now use the data and parity on the remaining disks to recreate the content of the failed disk, which reduces performance. In addition, while in degraded mode, the LUN is susceptible to the failure of the second disk. If a second disk in the LUN fails while in degraded mode, parity can no longer be used and all data on the LUN becomes inaccessible.

Figure 21 illustrates the distribution of user and parity data in a five-disk RAID 3 LUN. The the stripe segment size is 8 blocks, and the stripe size is 40 blocks (8 blocks times 5 disks). The disk block addresses in the stripe proceed sequentially from the first disk to the second, third, and fourth, then back to the first, and so on.

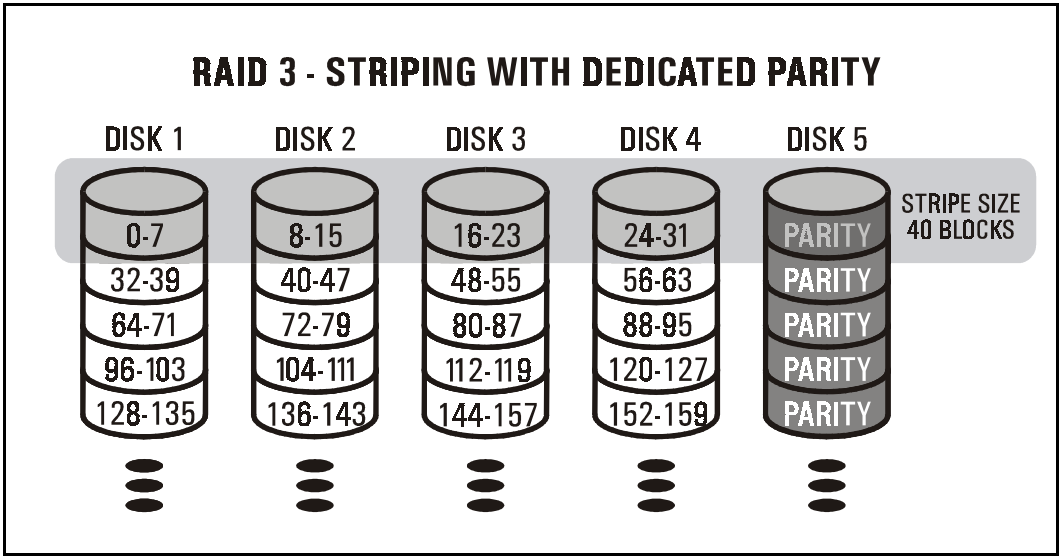



**Figure 21** RAID 3 LUN

RAID 3 works well for single-task applications using large block I/Os. It is not a good choice for transaction processing systems because the dedicated parity drive is a performance bottleneck. Whenever data is written to a data disk, a write must also be performed to the parity drive. On write operations, the parity disk can be written to four times as often as any other disk module in the group.

### RAID 5

RAID 5 uses parity to achieve data redundancy and disk striping to enhance performance. Data and parity information is distributed across all the disks in the RAID 5 LUN. A RAID 5 LUN consists of three or more disks. For highest availability, the disks in a RAID 5 LUN must be in different enclosures.

If a disk fails or becomes inaccessible, the disk array can dynamically reconstruct all user data from the data and parity information on the remaining disks. When a failed disk is replaced, the disk array automatically rebuilds the contents of the failed disk on the new disk. The rebuilt LUN contains an exact replica of the information it would have contained had the disk not failed.

Until a failed disk is replaced (or a rebuild on a global hot spare is completed), the LUN operates in degraded mode. The LUN must now use the data and parity on the remaining disks to recreate the content of the failed disk, which reduces performance. In addition, while in degraded mode, the LUN is susceptible to the failure of the second disk. If a second disk in the LUN fails while in degraded mode, parity can no longer be used and all data on the LUN becomes inaccessible.

Figure 22 illustrates the distribution of user and parity data in a five-disk RAID 5 LUN. The the stripe segment size is 8 blocks, and the stripe size is 40 blocks (8 blocks times 5 disks). The disk block addresses in the stripe proceed sequentially from the first disk to the second, third, fourth, and fifth, then back to the first, and so on.

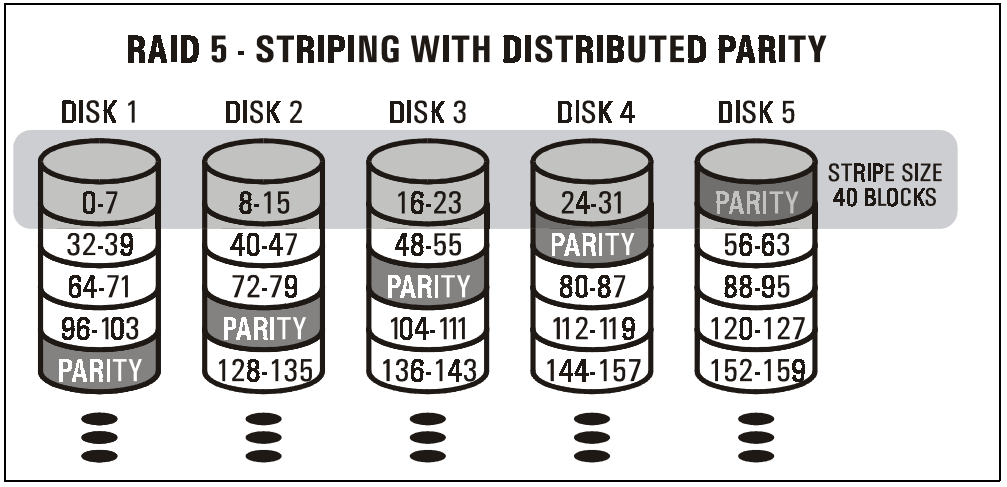



**Figure 22** RAID 5 LUN

With its individual access characteristics, RAID 5 provides high read throughput for small block-size requests (2 KB to 8 KB) by allowing simultaneous read operations from each disk in the LUN. During a write I/O, the disk array must perform four individual operations, which affects the write performance of a RAID 5 LUN. For each write, the disk array must perform the following steps:

1. Read the existing user data from the disks.
2. Read the corresponding parity information.
3. Write the new user data.
4. Calculate and write the new parity information.

Write caching can significantly improve the write performance of a RAID 5 LUN. RAID 5 is good for parallel processing (multi-tasking) applications and environments. The performance of a RAID 5 LUN is best when the maximum number of disks (six) is used.

## RAID Level Comparisons

To help you decide which RAID level to select for a LUN, the following tables compare the characteristics for the supported RAID levels. Where appropriate, the relative strengths and weakness of each RAID level are noted.

**Note** RAID 3 is supported on Windows NT and Windows 2000 only.

**Table 1** RAID Level Comparison: Data Redundancy Characteristics

| RAID Level | Disk Striping | Mirroring | Parity | Handle multiple disk failures? |
|---|---|---|---|---|
| RAID 0 | Yes | No | No | No. RAID 0 offers no data redundancy or protection against disk failure. RAID 0 should only be used for non-critical data. The failure of a single disk in a RAID 0 LUN will result in the loss of all data on the LUN. |
| RAID 1 | No | Yes | No | No |
| RAID 0/1 | Yes | Yes | No | Yes, providing both disks in a mirrored pair do not fail. |
| RAID 3 | Yes | No | Yes | No |
| RAID 5 | Yes | No | Yes | No |

**Table 2** RAID Level Comparison: Storage Efficiency Characteristics

| RAID Level | Storage Efficiency |
|---|---|
| RAID 0 | 100%. All disk space is use for data storage. |
| RAID 1 and 0/1 | 50%. All data is duplicated, requiring twice the disk storage for a given amount of data capacity. |
| RAID 3 and 5 | One disk's worth of capacity from each LUN is required to store parity data. As the number of disks in the LUN increases, so does the storage efficiency.<br><br>3-disk LUN: 66%<br>4-disk LUN: 75%<br>5-disk LUN: 80%<br>6-disk LUN: 83% |

**Table 3** RAID Level Comparison: Relative Performance Compared to an Individual Disk*

| LUN Configuration | Relative Read Performance for Large Sequential Access | Relative Write Performance for Large Sequential Access |
|---|---|---|
| RAID 0 | The read and write performance of a RAID 0 LUN increases as the multiple of the number of disks in the LUN. For example, a 4-disk RAID 0 LUN will achieve close to four times the performance of a single disk. | |
| RAID 1 mirrored pair | Up to 2.0 > than single disk | Equal to single disk |
| RAID 0/1 group with 10 disks | Up to 10.0 > than single disk | Up to 5.0 > than single disk |
| RAID 0/1 group with 6 disks | Up to 6.0 > than single disk | Up to 3.0 > than single disk |
| RAID 3 group with 5 disks | Up to 4.0 > than single disk | Up to 1.25 > than single disk |
| RAID-5 group with 5 disks | Up to 4.0 > than single disk | Up to 1.25 > than single disk |

* Compares the relative read and write performance for array configurations with the performance of a single stand-alone disk whose performance is 1.0. The read and write performance shown is the theoretical maximum performance relative to individual disk performance. The performance numbers are not based on read/write caching. With caching, the performance numbers for RAID 5 writes improve significantly.

**Table 4** RAID Level Comparison: General Performance Characteristics

| RAID Level | General Performance Characteristics |
|---|---|
| RAID 0 | – Simultaneous access to multiple disks increases I/O performance. In general, the greater the number of mirrored pairs, the greater the increase in performance. |
| RAID 1 | – A RAID 1 mirrored pair requires one I/O operation for a read and two I/O operations for a write, one to each disk in the pair.<br>– The disks in a RAID 1 mirrored pair are locked in synchronization, but the disk array can read data from the module whose read/write heads are the closest.<br>– RAID 1 read performance can be twice that of an individual disk. Write performance can be the same as that of an individual disk. |
| RAID 0/1 | – Simultaneous access to multiple mirrored pairs increases I/O performance. In general, the greater the number of mirrored pairs, the greater the increase in performance. |
| RAID 3 | – Provides high read throughput for large sequential I/Os.<br>– Write performance is limited by the need to perform four I/O operations per write request.<br>– Because some I/O operations occur simultaneously, performance depends on the number of disks in the LUN. Additional disks may improve performance.<br>– The I/O performance of RAID 5 benefits significantly from write caching. |
| RAID 5 | – Provides high read throughput for small block-size requests (2 KB to 8 KB).<br>– Write performance is limited by the need to perform four I/O operations per write request.<br>– Because some I/O operations occur simultaneously, performance depends on the number of disks in the LUN. Additional disks may improve performance.<br>– The I/O performance of RAID 5 benefits significantly from write caching. |

**Table 5** RAID Level Comparison: Application and I/O Pattern Performance Characteristics

| RAID level | Application and I/O Pattern Performance |
|---|---|
| RAID 0 | RAID 0 is a good choice in the following situations:<br>– Data protection is not critical. RAID 0 provides no data redundancy for protection against disk failure.<br>– Useful for scratch files or other temporary data whose loss will not seriously impact system operation.<br>– High performance is important. |
| RAID 1 | RAID 1 is a good choice in the following situations:<br>– Speed of write access is important.<br>– Write activity is heavy.<br>– Applications need logging or recordkeeping.<br>– Daily updates need to be stored to a database residing on a RAID 5 group. The database updates on the RAID 1 group can be copied to the RAID 5 group during off-peak hours. |
| RAID 0/1 | RAID 0/1 is a good choice in the following situations:<br>– Speed of write access is important.<br>– Write activity is heavy.<br>– Applications need logging or recordkeeping.<br>– Daily updates need to be stored to a database residing on a RAID 5 group. The database updates on the RAID 1 group can be copied to the RAID 5 group during off-peak hours. |
| RAID 3 | RAID 3 is a good choice in the following situations:<br>– Applications using I/O large sequential transfers of data, such as multimedia applications.<br>– Applications on which write operations are 33% or less of all I/O operations. |
| RAID 5 | RAID 5 is a good choice in the following situations:<br>– Multi-tasking applications using I/O transfers of different sizes.<br>– Database repositories or database servers on which write operations are 33% or less of all I/O operations.<br>– Multi-tasking applications requiring a large history database with a high read rate.<br>– Transaction processing is required. |

## Global Hot Spare Disks

A global hot spare disk is reserved for use as a replacement disk if a data disk fails. Their role is to provide hardware redundancy for the disks in the array. To achieve the highest level of availability, it is recommended that one global hot spare disk be created for each channel. A global hot spare can be used to replace any failed data disk within the array regardless of what channel it is on.

When a disk fails, the disk array automatically begins rebuilding the failed disk's data on an available global hot spare. When all the data has been rebuilt on the global hot spare, the LUN functions normally, using the global hot spare as a replacement for the failed disk. If a global hot spare is not available, data is still accessible using the redundant data maintained by the LUN.

When the failed disk is replaced, all data is copied from the former global hot spare onto the replacement disk. When the copy is complete, the former global hot spare is returned to the global hot spare disk group and is again available as protection against another disk failure.

If a failed disk is replaced while data is being rebuilt on the global hot spare, the rebuild process continues until complete. When all data is rebuilt on the global hot spare, it is then copied to the replacement disk.

Global hot spares are an essential component for maintaining data availability. A global hot spare reduces the risk of a second disk failure and restores the disk array's performance, which may be degraded while the LUN is forced to recreate data from parity. The use of multiple global hot spares may be desirable in environments where data availability is crucial. Multiple global hot spares ensure that data remains accessible even if multiple disks fail.

### Rebuilding Data

The rebuild process occurs any time a disk failure occurs. It uses the existing data and parity or mirror disk to rebuild the data that was on the failed disk. Because it is competing with host I/Os for disk array resources, a rebuild may affect disk array performance. The effect on performance is controlled by the rebuild priority settings. These settings determines how the disk array divides resources between the rebuild and host I/Os.

Settings that give a higher priority to the rebuild process will cause the rebuild to complete sooner, but at the expense of I/O performance. Lower rebuild priority settings favors host I/Os, which will maintain I/O performance but delay the completion of the rebuild.

The rebuild priority settings selected reflect the importance of performance versus data availability. The LUN being rebuilt is vulnerable to another disk failure while the rebuild is in progress. The longer the rebuild takes, the greater the chance of another disk failure.

The following sequence occurs following a disk failure and replacement. Figure 23 illustrates the process. A 5-disk RAID 5 LUN is used for this example.

1. Disk 3 in the RAID 5 LUN fails.
2. The disk array locates an available global hot spare and begins recreating on it the information that was on the failed disk. The data and parity on the remaining four disks in the LUN are used to recreate the information.
3. When the rebuild finishes, the global hot spare is part of the LUN, which fulfills the roll of disk 3.
4. When disk 3 is replaced, the disk array begins copying all the information from the former global hot spare to the replacement disk.
5. When copying completes, the LUN is restored to its original configuration. The former global hot spare is returned to the global hot spare disk group and is available to protect against another data disk failure.

**Note**

**Can a lower capacity disk serve as a hot spare for a larger disk?**
It is possible for a lower capacity disk to be used as a global hot spare when a larger disk fails. When a disk failure occurs, the disk array controller looks for a global hot spare that is large enough to store the data on the failed disk, not for a disk that matches the capacity of the failed disk. For example, if an 18 Gbyte disk fails but there is only 6 Gbytes of data stored on the disk, a 9 Gbyte global hot spare could be used.

Although this feature is available, it is recommended that you always select the largest disks in the array to serve as global hot spares. This will ensure that any disk in the array is protected, regardless of how much data is stored on it.

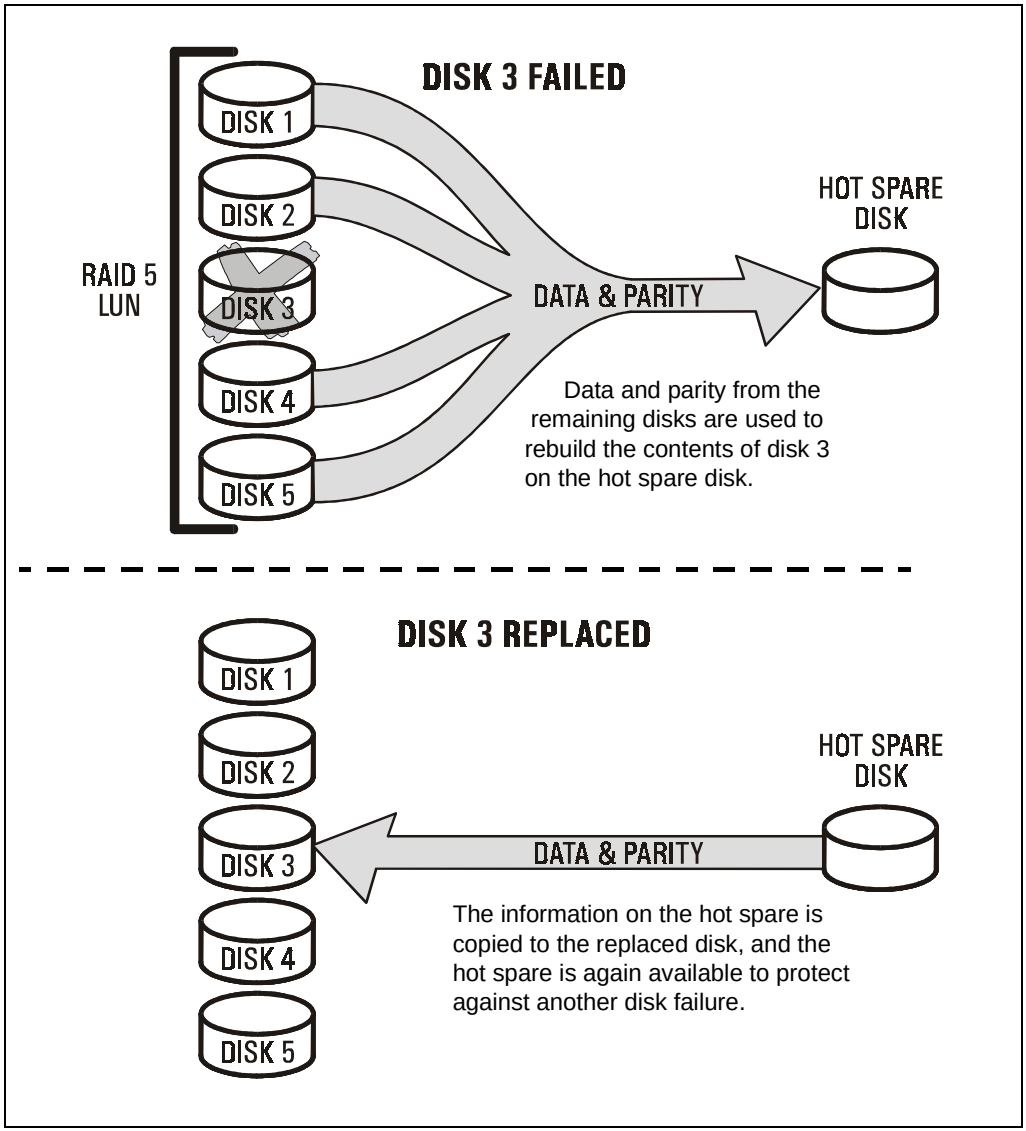

DISK 3 FAILED
DISK 1
DISK 2
DISK 3
DISK 4
DISK 5
RAID 5
LUN
HOT SPARE
DISK
DATA & PARITY
Data and parity from the remaining disks are used to rebuild the contents of disk 3 on the hot spare disk.
DISK 3 REPLACED
DISK 1
DISK 2
DISK 3
DISK 4
DISK 5
HOT SPARE
DISK
DATA & PARITY
The information on the hot spare is copied to the replaced disk, and the hot spare is again available to protect against another disk failure.

**Figure 23** Rebuild Process on a RAID 5 LUN (or Volume Group)

## Primary and Alternate I/O Paths

There are two I/O paths to each LUN on the disk array - one through controller A and one through controller B. Logical Volume Manager (LVM) is used to establish the primary path and the alternate path to a LUN. The primary path becomes the path for all host I/Os to that LUN.

If a failure occurs in the primary path, LVM automatically switches to the alternate path to access the LUN. The first time an I/O is performed to the LUN using the alternate path, the disk array switches ownership of the LUN to the controller on the alternate path. Once the problem with the primary path is corrected, ownership of the LUN should be switched back to the original I/O path to maintain proper load balancing.

The primary path established using LVM defines the owning controller for the LUN. This may override the controller ownership defined when the LUN was bound. For example, if controller A was identified as the owning controller when the LUN was bound, and LVM subsequently established the primary path to the LUN through controller B, controller B becomes the owning controller.

# Capacity Management Features

The disk array uses a number of features to manage its disk capacity efficiently. The use of LUNs allow you to divide the total disk capacity into smaller, more flexible partitions. Caching improves disk array performance by using controller RAM to temporarily store data during I/Os.

**Note** **Differences in HP-UX and Windows NT/2000 Capacity Management**
Capacity management on Windows NT and Windows 2000 offers some unique features. Refer to the *HP Storage Manager 60 Introduction Guide* for information on Windows-specific features. Some of the terms used in the *HP Storage Manager 60* software differ from those used here. These terms are also listed in the *HP Storage Manager 60 Introduction Guide*.

## LUNs

The capacity of the disk array can be divided into entities called LUNs. Individual disks are grouped together to form a LUN. Functionally, each LUN appears to the host operating system as an individual disk drive.

Although the LUN appears to the host as an individual disk, the use of multiple disks offers advantages of increased data availability and performance. Data availability is enhanced by using redundant data stored on a separate disk from the original data. The use of multiple disks increases performance by allowing simultaneous access to several disks when reading and writing data.

## Disk Groups

A disk group is a collection of individual disks that share a common role in disk array operation. All disks on the disk array become a member of one of the following disk groups:

- LUN group – Each LUN on the disk array has its own disk group. When a disk is included as part of a newly created LUN, the disk becomes a member of the associated disk group. There can be only one LUN in each LUN disk group.

- Hot spare group – All disks assigned the role of global hot spare become members of this group. Up to six disks (one for each channel) can be assigned as global hot spares.
- Unassigned group – Any disk that is neither part of a LUN nor a global hot spare is considered unassigned and becomes a member of this group. Unassigned disks can be used to create a LUN or can be used as global hot spares. Unassigned disks do not contribute to the capacity of the disk array.

## Disk Array Caching

Disk caching is the technique of storing data temporarily in RAM while performing I/Os to the disk array. Using RAM as a temporary storage medium can significantly improve the response time for many types of I/O operations. From the host's perspective the data transfer is complete, even if the disk media was not involved in the transaction. Both write caching and read caching are always enabled.

Caching enhances disk array I/O performance in two ways:

| | |
|---|---|
| Read I/O | If a read I/O requests data that is already in read cache, the disk array services the request from cache RAM, thus avoiding the much slower process of accessing a disk for the data. A pre-fetch capability enables the disk array to anticipate needed data (for example, on a file transfer) and read it from disk into the read cache, which helps significantly with sequential read I/Os. |
| Write I/O | During a write I/O, the disk array writes the requested data into write cache. Rather than writing the modified data back to the disk immediately, the disk array keeps it in cache and informs the host that the write is complete. If another I/O affects the same data, the disk array can update it directly in cache, avoiding another disk write. Data is flushed to disk at regular intervals (10 seconds) ro when the cache flush threshold is reached. |

Write cache is mirrored between the two disk array controllers. Each controller maintains an exact image of the write cache on the other controller. If a controller fails, its write cache content is flushed to the disk by the remaining controller. Because write cache is mirrored, the operational controller automatically disables write caching until the failed controller is replaced. After it is replaced, the operational controller automatically re-enables write caching. Mirroring effectively reduces the size of available cache by half. A

controller with 256 Mbytes of cache will use half of the memory to mirror the other controller, leaving only 128 Mbytes for its own cache.

The write cache contents cannot be flushed when both controllers are removed from the disk array simultaneously. In this case the write cache image is lost and data integrity on the disk array is compromised. To avoid this problem, *never* remove both controllers from the disk array simultaneously.

In the event of an unexpected disk array shutdown or loss of power, the BBU provides power to cache memory to maintain the cache for 120 hours (5 days).

## Dynamic Capacity Expansion

If slots are available in the disk enclosures, you can increase the capacity of the disk array without disrupting operation. By simply adding new disks to the array and then creating a new LUN the capacity can be expanded. See "Adding Capacity to the Disk Array" on page 254 for more information on adding disks and other ways of increasing disk array capacity.

# 2 TOPOLOGY AND ARRAY PLANNING



Overview . . . . . 70

Array Design Considerations . . . . . 71

Recommended Disk Array Configurations . . . . . 77

Topologies for HP-UX. . . . . 102

Topologies for Windows NT and Windows 2000 . . . . . 131

# Overview

This chapter provides information to assist you in configuring the Disk Array FC60 to meet your specific storage needs. Factors to be considered when configuring the disk array include high availability requirements, performance, storage capacity, and future expandability. This chapter discusses configuration features of the Disk Array FC60 as it relates to these requirements. In addition it provides information on system topologies following the array configuration.

**Note** Planning the disk array configuration is typically done before the disk array hardware is received. This information is then used during installation of the disk array to create the desired system configuration. After the disk array is installed, the information in this chapter will help you verify your configuration.

# Array Design Considerations



The Disk Array FC60 provides the versatility to meet varying application storage needs. To meet a specific application need, the array should be configured to optimize the features most important for the application. Array features include:

- High availability
- Performance (high I/O transfer rates)
- Storage capacity (optimize for lowest cost/Mbyte)
- Scalability

Optimizing array operation for a specific feature is managed by configuring various array installation and operation options. These options include: hardware configuration, RAID level and LUN creation, the number of SC10 disk enclosures used, and system support software. It should be noted that optimizing the disk array for one feature may compromise another. For example, optimizing for maximum performance may increase the cost per megabyte of storeage.

## Array Hardware Configuration

Array configuration options that affect high availability, performance, and storage capacity include Fibre Channel connections, disk enclosure bus configuration, and internal Ultra2 SCSI channels. This information is presented first, because it is the basis for some of the array configuration planning.

### Fibre Channel Connection

If the controller enclosure has both controller modules installed, dual Fibre Channels can be connected to the controller enclosure (if only one controller module is installed, only one Fibre Channel cable can be connected). Using dual Fibre Channel-AL connections increases the data throughput and provides for higher data availability.

### Ultra2 SCSI Channel Operations

The disk array controller enclosure provides six Ultra2 SCSI channel connections for up to six disk enclosures. Six separate SCSI channels provide configuration flexibility. Disk

enclosures can be added incrementally (up to six) as storage requirements grow. Multiple SCSI channels also increase data throughput. This increased data throughput occurs as a result of the controller's ability to transfer data simultaneously over multiple data paths (channels). The more channels used, the faster the data throughput.

### Disk Enclosure Bus Configuration

The disk enclosure can connect to either one or two SCSI channels, depending on its bus configuration. Disk enclosure design allows the backplane bus to be split from a single bus into two separate SCSI busses. When the backplane is operating as a single bus it is referred to as full-bus mode; when the bus is split into two separate buses it is referred to as split-bus mode. When in full-bus mode, one SCSI channel connects to all of the ten disk modules. If the enclosure is configured for split-bus mode, five disk modules are connected to each of the two separate buses and a separate SCSI channel connects to each of the BCCs. See "Installation" on page 143 for more information.

When using split-bus mode, the maximum number of disk enclosures is limited to three (each disk enclosure uses two channel connections). If the storage capacity needs to be increased by adding more disk enclosures, the array will need to be reconfigured. Reconfiguring the array requires shutting down the host system, powering down the array, installing additional disk enclosures, and reconfiguring and recabling all enclosures.

## RAID, LUNs, and Global Hot Spares

In addition to the above hardware configuration considerations, the RAID level and LUN structure has considerable impact on high availability, performance, and storage capacity. For information on how RAID level, LUNs, and hot spares affect the performance of a disk array, see "Disk Array High Availability Features" on page 47.

## High Availability

If your application requires high availability, you should implement the options discussed here.The Disk Array FC60 is fully qualified to run MC/ServiceGuard and MC/LockManager. To work in these environments, a high availability configuration must be used. To configure

the array for high availability, there must be no single points of failure. This means that the configuration must have at least these minimum characteristics:

- Two controllers connected to separate Fibre Channel loops (using separate Fibre Channel host I/O adaptors)
- Two disk enclosures (minimum)
- Eight disk modules, four in each disk enclosure (minimum)
- LUNs that use only one disk per disk enclosure.

With its dual controllers, the Disk Array FC60 provides two independent I/O paths to the data stored on the array. Data is transmitted from the array controllers to the disks through up to six Ultra2 SCSI channels connected to the disk enclosures. Any of several RAID levels (1, 0/1, or 5) can be selected, however, for the greatest high availability, It is recommended that RAID level 1 be used for optimum high availability.

**Note** The Disk Array FC60 is designed to operate with either one or two controller modules, however for high availability it is highly recommended that two controller modules be installed.

## Performance

The maximum aggregate performance that can be sustained by the disk array is approximately 170 megabytes per second (using dual Fibre Channel-AL connections). This performance can be achieved by configuring at least four disk modules per Ultra2 SCSI bus utilizing the six Ultra2 SCSI channels. This can be accomplished in two ways. One way is to configure six disk enclosures, one each per Ultra2 SCSI channel (disk enclosure full-bus mode). In each of these enclosures, configure at least four disk modules.

Adding more disk modules to each of these disk enclosure will increase storage capacity, but will not appreciably increase the sequential thoughput. Additional capacity may be a worthwhile addition, since in many computing environments, capacity, not access speed, is the limiting factor.

Another way to configure for maximum performance is to connect three disk enclosures to the controller enclosure and configure these enclosures for split-bus operation. Then connect an Ultra2 SCSI channel to each split bus (two channels per disk enclosure). Each

of the buses must be configured with at least four disk modules (eight disk modules per disk enclosure). This configuration also offers full sequential performance and is more economical to implement.

To scale up sequential transfer performance from the host, configure additional disk arrays. This will increase the total I/O bandwidth available to the server.

Performance can also be measured by the number of I/O operations per second a system can perform. I/Os per second are important in OLTP (on-line transaction processing) applications. To maximize I/Os per second, configure the maximum number of disk modules. For the same capacity, you may elect to use a larger number of 9.1 Gbyte disk modules instead of a smaller number of higher capacity disk modules to obtain optimal I/Os per second.

**Note** For the maximum I/Os per second, configure RAID 0/1 and the maximum number of disk modules.

## Storage Capacity

For configurations where maximum storage capacity at minimum cost is a requirement, consider configuring the disk array in RAID 5 (using the maximum number of data drives per parity drives) and only supplying one or two hot spare drives per disk array. Also, purchase the lowest cost/Mbyte drive available (typically the largest capacity drives available at the time of purchase). This configuration allows the maximum volume of storage at the lowest cost. Disk arrays configured in this way will need to be carefully monitored to make sure that failed disks are promptly replaced.



## Expanding Storage Capacity

The disk array is designed to meet a range of capacity requirements. It can be configured with up to six disk enclosures, and from 4 to 60 disk modules. The disk array can be purchased with enough capacity to meet your current storage requirements. As your storage needs grow, you can easily add more capacity to the Disk Array FC60.

There are several ways to increase the storage capacity of the disk array. It can be increased by replacing existing smaller capacity disk modules with larger capacity disk modules, by adding more disk modules to the disk enclosures, or by adding additional disk enclosures.

The best method for expansion is to install all six disk enclosures at initial installation (full-bus configuration). Then, install only the required capacity (number of disk modules), leaving empty disk slots for future expansion. Additional disk modules can be installed as the requirement for additional capacity grows. This method allows for greater flexibility in LUN creation and does not require that the system be shut down for expansion. Adding disk enclosures to the array after initial installation requires that the system be shut down during the installation. All the remaining configuration expansion methods require that the data to the array be suspended (the system shut down) to add additional disk enclosure storage to the array.

**Note** For maximum performance, all six SCSI channels from the controller enclosure must be connected with a minimum of two disk modules per channel (disk enclosure configured for full bus mode).

To expand a split-bus configuration (adding more disk enclosures) will require shutting the host system down. If the initial installation included only one or two disk enclosures, then

another, two or one disk enclosures, respectively, can be added by using split-bus mode. However, if you are adding up to four, five, or six enclosure, the enclosures configuration will need to be switched from split-bus to full-bus (refer to “Disk Enclosure Bus Configuration” section, earlier in this chapter, for additional information).

**Note** Typically adding only one enclosure does not provide any options for creating LUNs. It is best to expand the array with at least a minimum of two disk enclosures at a time. Two additional drives could then be configured as RAID 1 or 0/1.

Installing a five disk enclosure array limits expansion to six disk enclosures. Adding one additional enclosure does not provide any versatility for creating LUNs (unless all data is removed and the LUNs are rebuilt).

If the initial installation uses split-bus disk enclosure (split-bus accepts three disk enclosures maximum) and expansion requires adding four or more enclosures, the existing disk enclosures will need to be reconfigured for full-bus mode and the additional number of enclosures installed into the rack and required number of full-bus enclosures. This expansion requires disk enclosures be recabled using full bus enclosures. As in all cases of adding disk enclosures to the array, the system has to be shut down for the expansion. Determine the RAID level and how the LUNs will be created for the expansion

If the initial installation consisted of one or more full-bus configured disk enclosures, then additional full-bus configured disk enclosures can be added to the array. The system should be shut down for the addition of the enclosures. Determine the RAID level and how the LUNs will be created for the additional storage.

To scale up sequential performance, first make sure that the configuration includes both controllers modules. Maximum sequential transfer performance will be reached with approximately 20 disk modules simultaneously transferring data. To achieve additional sequential transfer performance, you will need to add a second disk array and more disk modules.

To increase I/Os per second performance, add disk modules. Transaction performance is directly related to the number of disk modules installed in the disk array.

# Recommended Disk Array Configurations

This section presents recommended configurations for disk arrays using one to six disk enclosures. Configurations are provided for achieving high availability/high performance, and maximum capacity. The configuration recommended by Hewlett-Packard is the high availability/ high performance configuration, which is used for factory assembled disk arrays (A5277AZ). The configurations identify the number of disk enclosures, cable connections, disk enclosure bus modes, RAID level, and LUN structure.

Most of the configurations offer the highest level of availability, which means they are capable of surviving the failure of a single disk (provided LUNs are created with one disk module per disk enclosure), SCSI channel, disk enclosure, or controller module. The only configurations that do not offer the highest level of availability are the single disk enclosure configuration, and the two enclosure high capacity configuration. These configurations cannot survive the failure of an entire disk enclosure, so they should not be used in environments where high availability is critical.

The configurations list maximum disk capacity and usable disk capacity (with ten disk modules installed). Configurations based on RAID 1 have less usable disk capacity than RAID 5, but I/O performance is optimized when using RAID 1. Although the recommended configurations presented here all contain ten disk modules, a disk enclosure can contain four, eight, or ten disk modules.

**Note** The terms “LUN” and “volume group” are used interchangeably in the text and figures in this section.

## Configuration Considerations

The following factors should be considered when using any of the recommended configurations.

- **Multiple Hosts** - A single host system is shown, but configurations can be adapted to create multi-host, high availability systems. For more information on using multiple hosts, see "Topologies for HP-UX" on page 102 or "Topologies for Windows NT and Windows 2000" on page 131.

- **Global hot spares** - although none of the configurations use global hot spares, their use is recommended to achieve maximum protection against disk failure. For more information, see "Global Hot Spare Disks" on page 61.
- **Split bus operation** - With three or fewer disk enclosures, increased performance can be achieved by operating the disk enclosures in split bus mode, which increases the number of SCSI busses available for data transfer. However, operating the disk enclosures in split bus mode may make if more difficult to expand the capacity of the array. In a split bus configuration, it may be necessary to take down the host, back up data, and rebind LUNs when adding disk enclosures. If you anticipate the need to expand your disk array, you may want to consider selecting a configuration that uses more enclosures operating in full bus mode. In addition to simplifying expansion, this type of configuration also gives you greater flexibility when creating LUNs.
- **Segment size -** the recommended segment size is 16 Kbyte for RAID 5, and 64 Kbyte for RAID 1 and RAID 0/1.
- **Maximum LUNs** - A maximum of 30 LUNs can be configured on the disk array.

## One Disk Enclosure Configuration

**Note** A single disk enclosure configuration is not recommended for environments where high availability is critical. For optimum high availability, at least two disk enclosures are required. This protect against the failure of a single disk enclosure.

- **Hardware Configuration**
  - Two disk array controllers connected directly to host Fibre Channel adapters
  - One disk enclosure with ten 73 GBytes disk modules
  - Disk enclosure configured for split-bus mode (two SCSI channels)
- **LUN Configuration**
  - Five RAID 1 LUNs, each comprising two disks (1+1)
  - Each disk in a LUN is on a separate SCSI channel

- **Data Availability**
  - Not recommended for maximum high availability.
  - Handles a single disk failure, single BCC failure, a single channel failure, or a single controller failure
  - Expansion requires powering down the disk array, removing terminators and/or cables from the enclosures, and cabling additional disk enclosures.
- **Disk Capacity**
  - Maximum capacity 730 GBytes
  - Usable capacity 365 GBytes

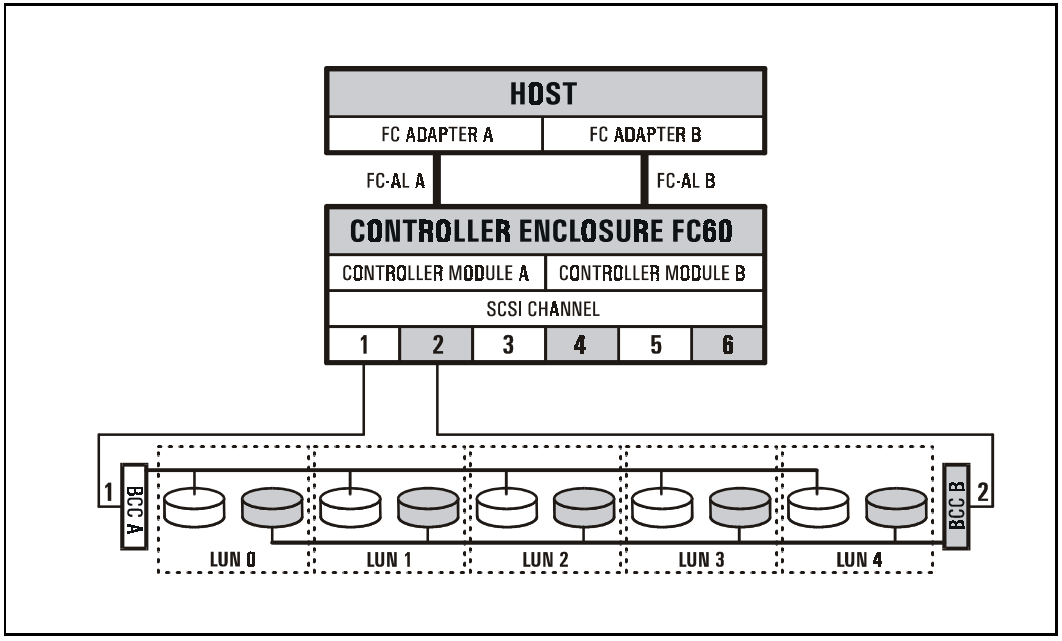



**Figure 24** One Disk Enclosure Array Configuration

## Two Disk Enclosure Configurations

### High Availability/ High Performance

- **Hardware Configuration**
  - Two disk array controllers connected directly to host Fibre Channel adapters
  - Two disk enclosures with ten 73 GByte disk modules (20 disks total)
  - Disk enclosures configured for split-bus mode (two SCSI channels per enclosure)
- **LUN Configuration**
  - Ten RAID 1 LUNs, each comprising two disks (1+1)
  - Each disk in a LUN is in a separate enclosure
- **High Availability**
  - Handles a single disk failure, BCC failure, single channel failure, or a single controller failure.
  - Expansion requires powering down the disk array, removing terminators and/or cables from the enclosures, and cabling additional disk enclosures.
- **Disk Capacity**
  - Maximum capacity 1460 GBytes
  - Usable capacity 730 GBytes

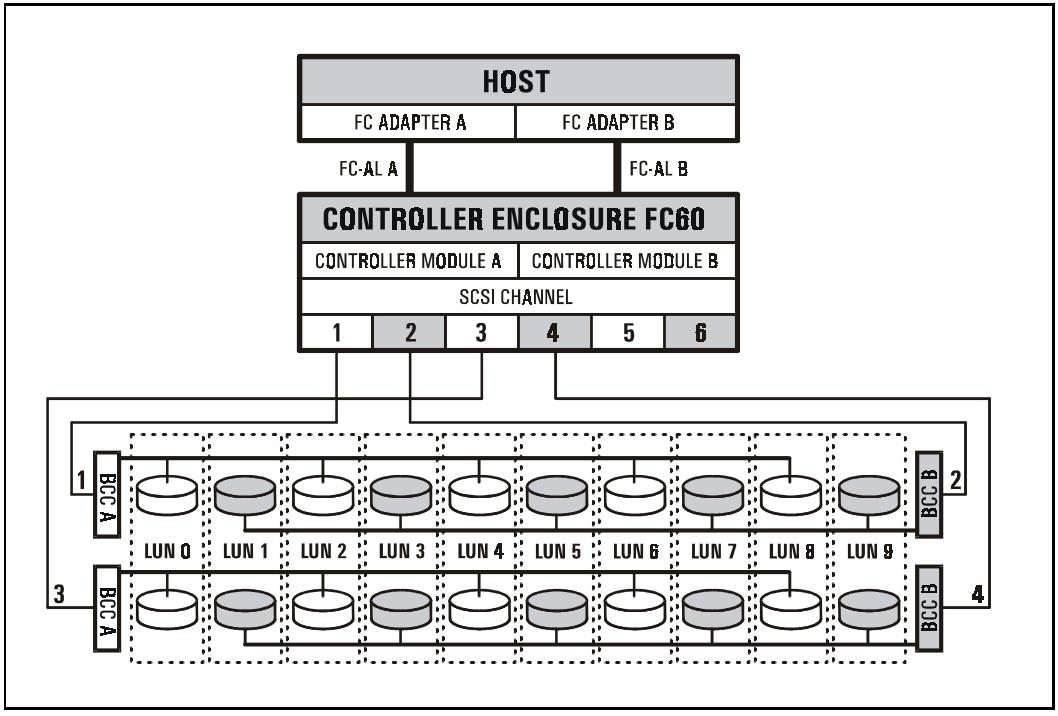



**Figure 25** Two Disk Enclosure High Availability/ High Performance Configuration

## Maximum Capacity

**Note** This configuration is not recommended for environments where high availability is critical. To achieve high availability each disk in a LUN should be in a different disk enclosure. This configuration does not achieve that level of protection.

- **Hardware Configuration**
  - Two disk array controllers connected directly to host Fibre Channel adapters
  - Two disk enclosures with ten 73 GByte disks each (20 disks total)
  - Disk enclosures configured for split-bus mode (two SCSI channels per enclosure)
- **LUN Configuration**
  - Five RAID 5 LUNs, each comprising four disks (3 data + 1 parity)
  - Each disk in a LUN is on a separate SCSI bus.
- **High Availability**
  - Handles the failure of a single disk, single controller, or a single channel
  - Does not handle a disk enclosure failure, consequently this configuration is NOT recommended for critical high availability installations.
  - Expansion requires powering down the disk array, removing terminators and/or cables from the enclosures, and cabling additional disk enclosures.
- **Disk Capacity**
  - Maximum capacity 1460 GBytes
  - Usable capacity 1095 GBytes

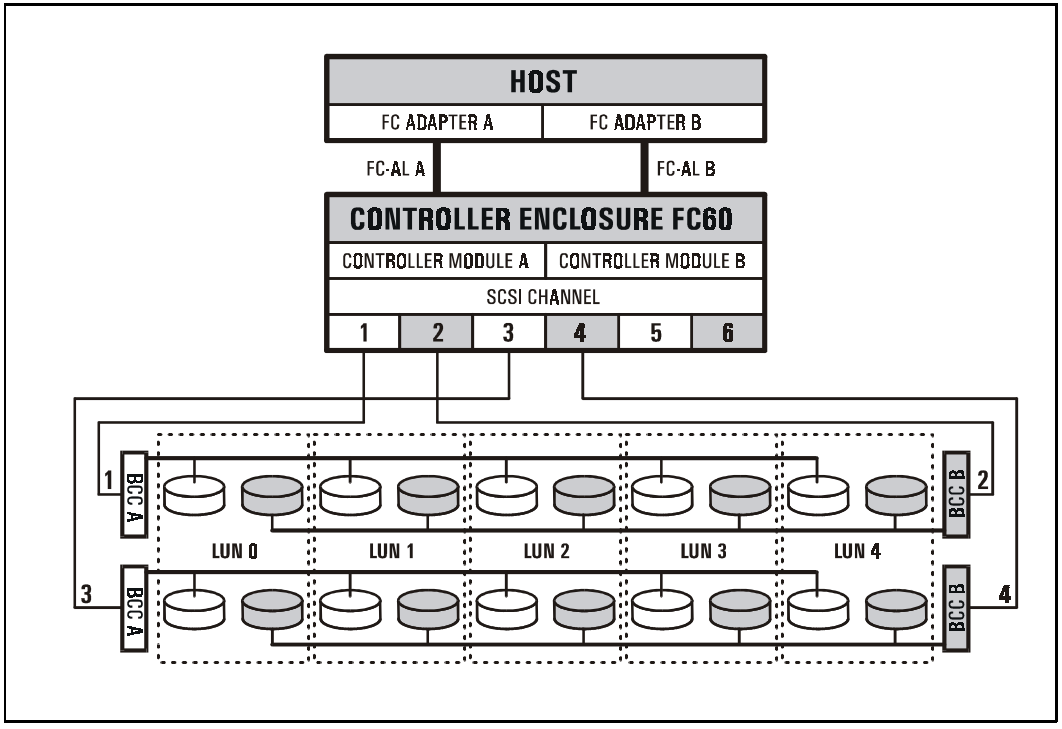



**Figure 26** Two Disk Enclosure Maximum Capacity Configuration

## Three Disk Enclosure Configurations

### High Availability/ High Performance

- **Hardware Configuration**
  - Two disk array controllers connected directly to host Fibre Channel adapters
  - Three disk enclosures with ten 73 GByte disks each (30 disks total)
  - Disk enclosures configured for split-bus mode (two SCSI channels per enclosure)
- **LUN Configuration**
  - 15 RAID 1 LUNs, each comprising two disks (1+1)
  - Each disk in a LUN is in a separate enclosure
- **High Availability**
  - Handles a single disk failure, a single controller, a single channel, single BCC, or a single disk enclosure failure
  - Expansion requires powering down the disk array, recabling the array to a full bus configuration, rebinding the LUNs, and restoring all data
- **Disk Capacity**
  - Maximum capacity 2190 GBytes
  - Usable capacity 1095 GBytes

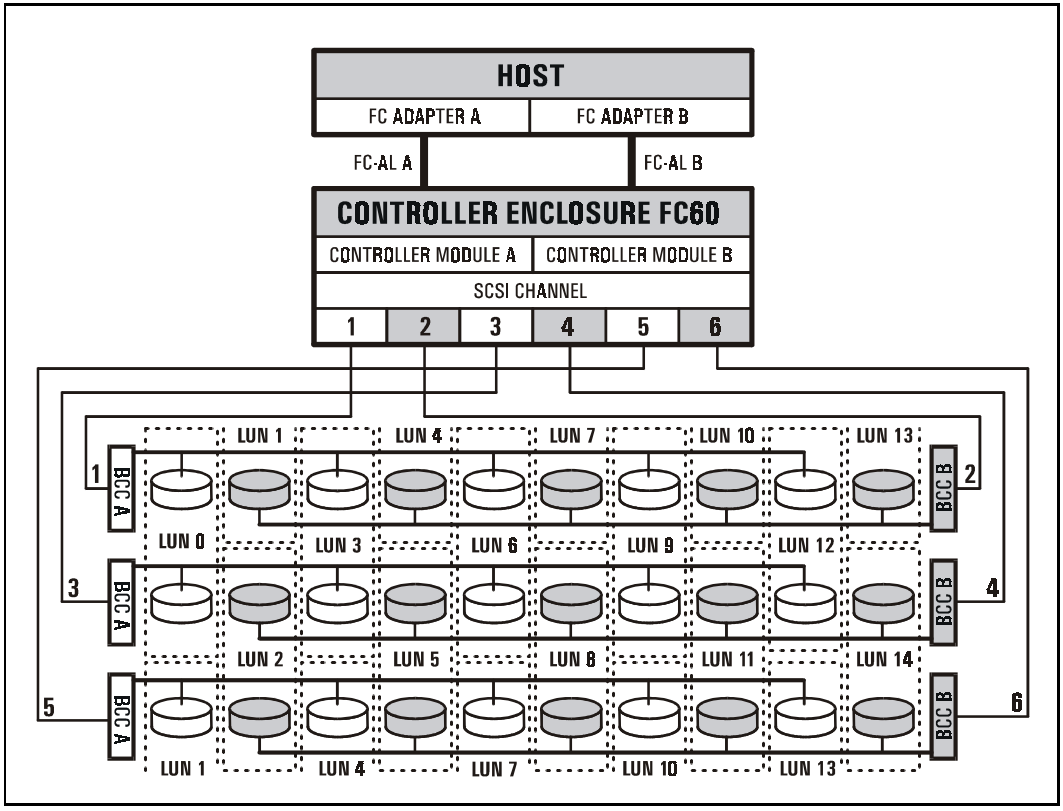



**Figure 27** Three Disk Enclosure High Availability/ High Performance Configuration

## Maximum Capacity

- **Hardware Configuration**
  - Two disk array controllers connected directly to host Fibre Channel adapters
  - Three disk enclosures with ten 73 GByte disks each (30 disks total)
  - Disk enclosures configured for split-bus mode (two SCSI channels per enclosure)
- **LUN Configuration**
  - Ten RAID 5 LUNs, each comprising three disks (2 data + 1 parity).
  - Each disk in a LUN is in a separate enclosure.
- **High Availability**
  - Handles a single disk failure, single controller, single channel, single BCC, or a single disk enclosure failure
  - Expansion requires powering down the disk array, recabling the array to a full bus configuration, rebinding the LUNs, and restoring all data
- **Disk Capacity**
  - Maximum capacity 2190 GBytes
  - Usable capacity 1460 GBytes

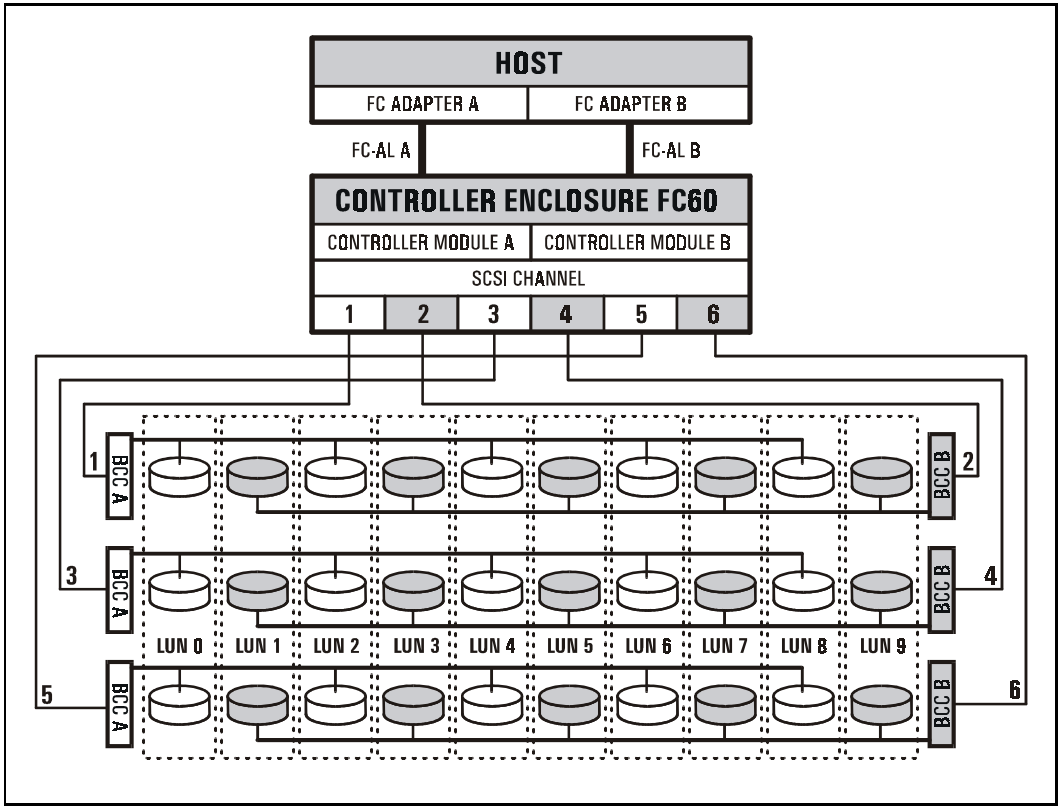



**Figure 28** Three Disk Enclosure Maximum Capacity Configuration

## Four Disk Enclosure Configurations

### High Availability/High Performance

- **Hardware Configuration**
  - Two disk array controllers connected directly to host Fibre Channel adapters
  - Four disk enclosures with ten 73 GByte disks each (40 disks total)
  - Disk enclosures configured for full-bus mode (one SCSI channel per enclosure)
- **LUN Configuration**
  - Ten RAID 0/1 LUNs, each comprising four disks (2+2)
  - Each disk in a LUN is in a separate enclosure.
- **High Availability**
  - Handles a single disk failure, single disk enclosure/BCC failure, single channel failure, or a single controller failure
  - Expansion requires powering down the disk array and adding additional disk enclosures and cables
- **Disk Capacity**
  - Maximum capacity 2920 GBytes
  - Usable capacity 1460 GBytes

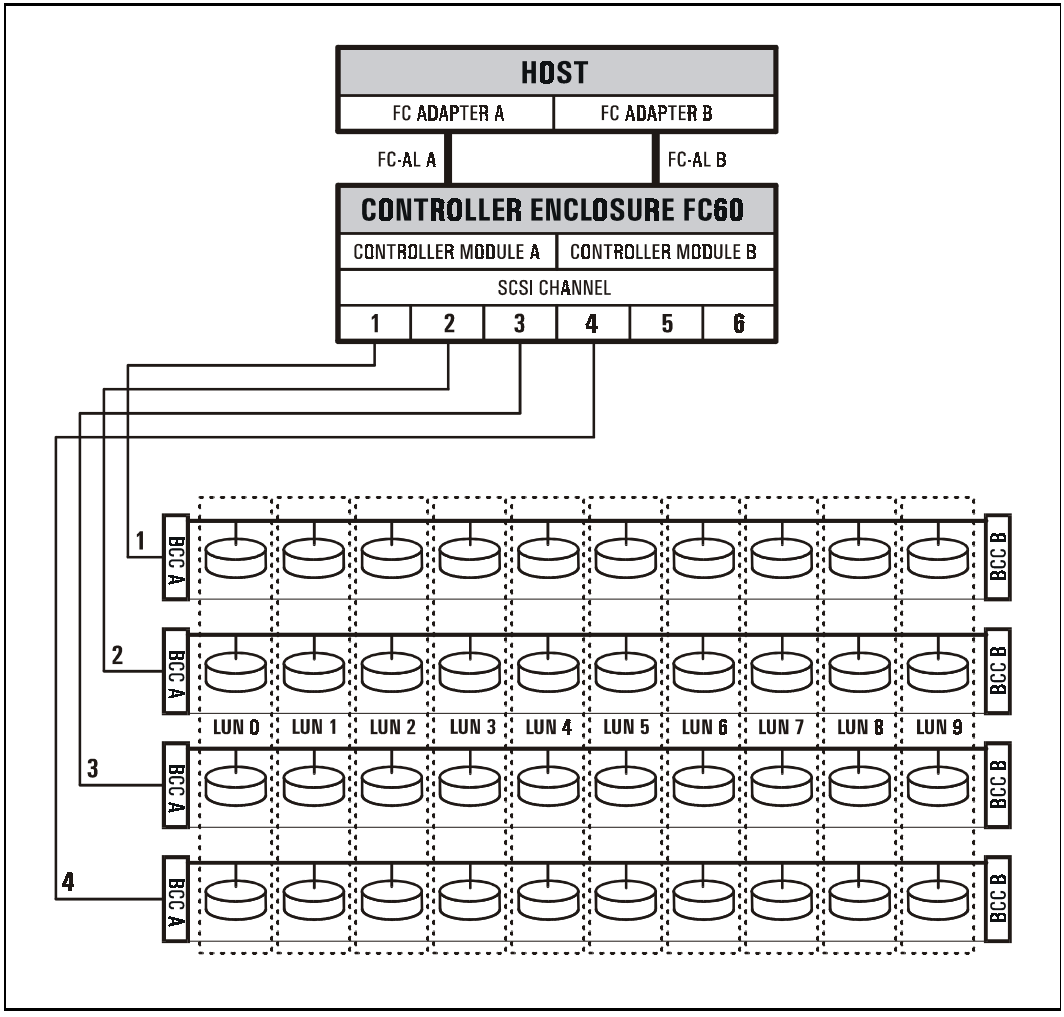



**Figure 29** Four Disk Enclosure High Availability/High Performance Configuration

## Maximum Capacity

- **Hardware Configuration**
  - Two disk array controllers connected directly to host Fibre Channel adapters
  - Four disk enclosures with ten 73 GByte disks each (40 disks total)
  - Disk enclosures configured for full-bus mode (one SCSI channel per enclosure)
- **LUN Configuration**
  - Ten RAID 5 LUNs, each comprising four disks (3 data + 1 parity)
  - Each disk in a LUN is in a separate enclosure.
- **High Availability**
  - Handles a single disk failure, single controller, single channel, single BCC, or a single disk enclosure failure
  - Expansion requires powering down the disk array and adding additional disk enclosures and cabling
- **Disk Capacity**
  - Maximum capacity 2920 GBytes
  - Usable capacity 2190 GBytes

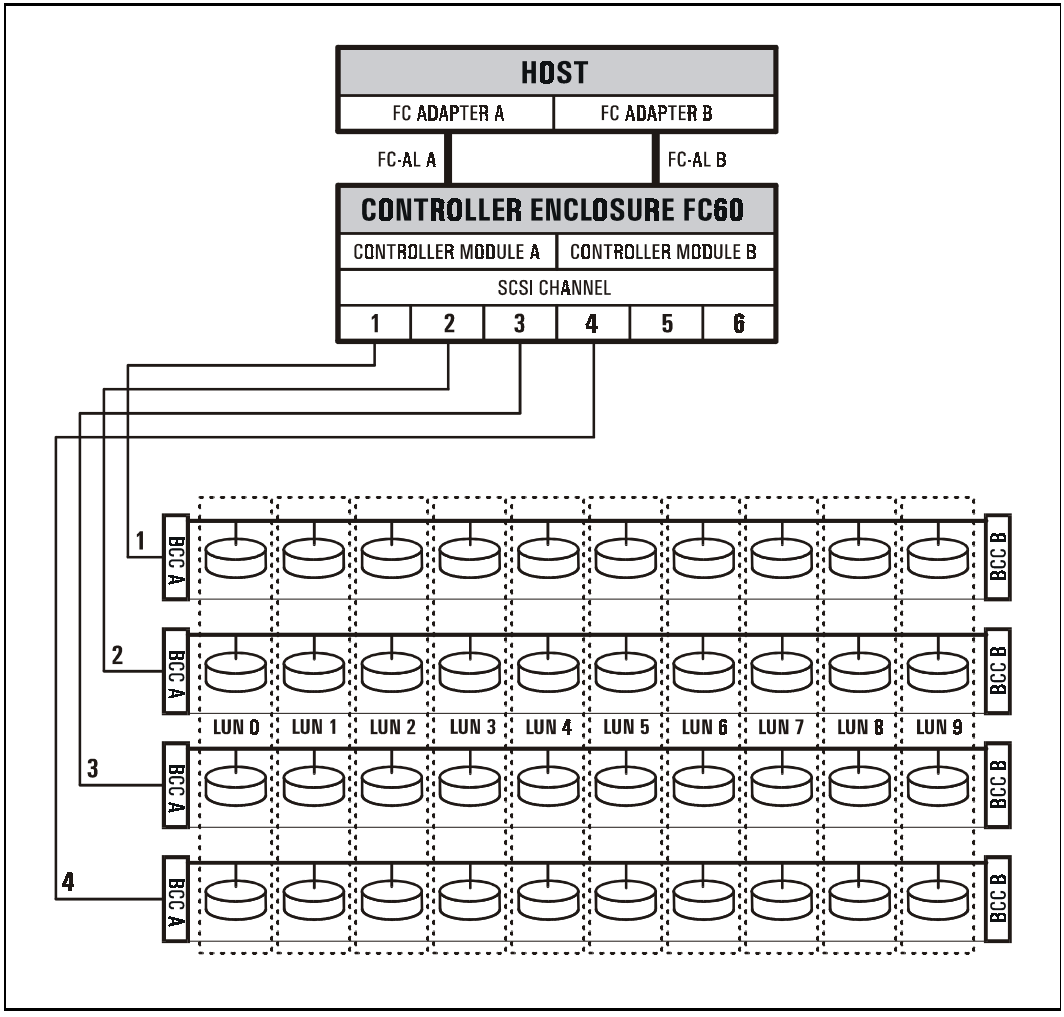



**Figure 30** Four Disk Enclosure Maximum Capacity Configuration

## Five Disk Enclosure Configurations

### High Availability/High Performance

- **Hardware Configuration**
  - Two disk array controllers connected directly to host Fibre Channel adapters
  - Five disk enclosures with ten 73 GByte disks each (50 disks total)
  - Disk enclosures configured for full-bus mode (one SCSI channel per enclosure)
- **LUN Configuration**
  - Ten RAID 0/1 LUNs, each comprising four disks (2+2)
  - Five RAID 1 LUNs, each comprising two disks (1+1)
  - Each disk in a LUN is in a separate enclosure.
- High Availability
  - Handles a single disk failure, BCC failure, single channel failure, or a single controller failure
  - Expansion requires powering down the disk array and adding an additional disk enclosure and cabling
- **Disk Capacity**
  - Maximum capacity 3650 GBytes
  - Usable capacity 1825 GBytes

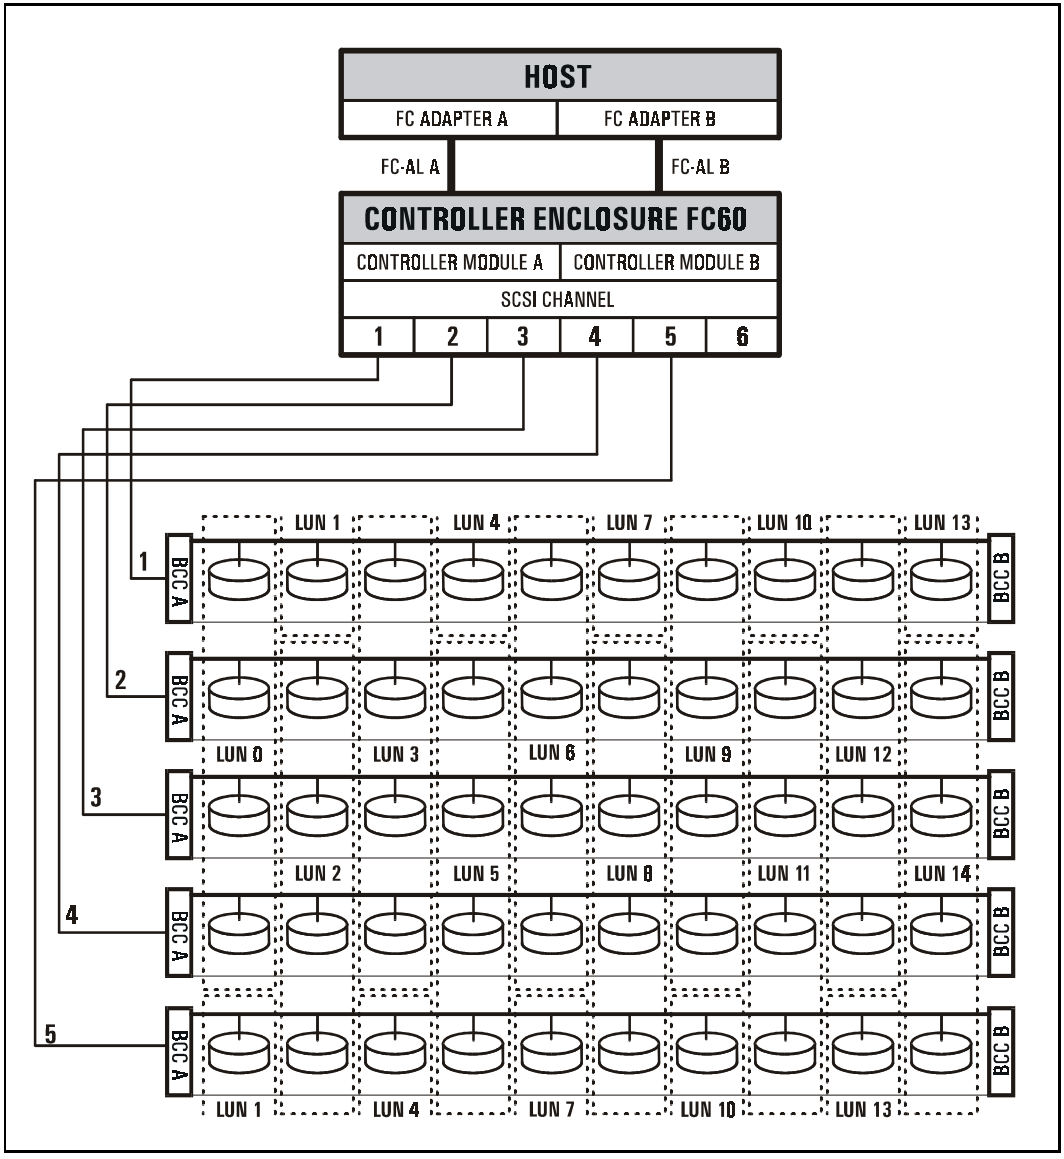



**Figure 31** Five Disk Enclosure High Availability/High Performance Configuration

## Maximum Capacity

- **Hardware Configuration**
  - Two disk array controllers connected directly to host Fibre Channel adapters
  - Five disk enclosures with ten 73 GByte disks each (50 disks total)
  - Disk enclosures configured for full-bus mode (one SCSI channel per enclosure)
- **LUN Configuration**
  - Ten RAID 5 LUNs, each comprising five disks (4 data + 1 parity)
  - Each disk in a LUN is in a separate enclosure.
- **High Availability**
  - Handles a single disk failure, single disk enclosure/BCC failure, single channel failure, or a single controller failure
  - Expansion requires powering down the disk array, and adding an additional disk enclosure and cabling
- **Disk Capacity**
  - Maximum capacity 3650 GBytes
  - Usable capacity 2920 GBytes

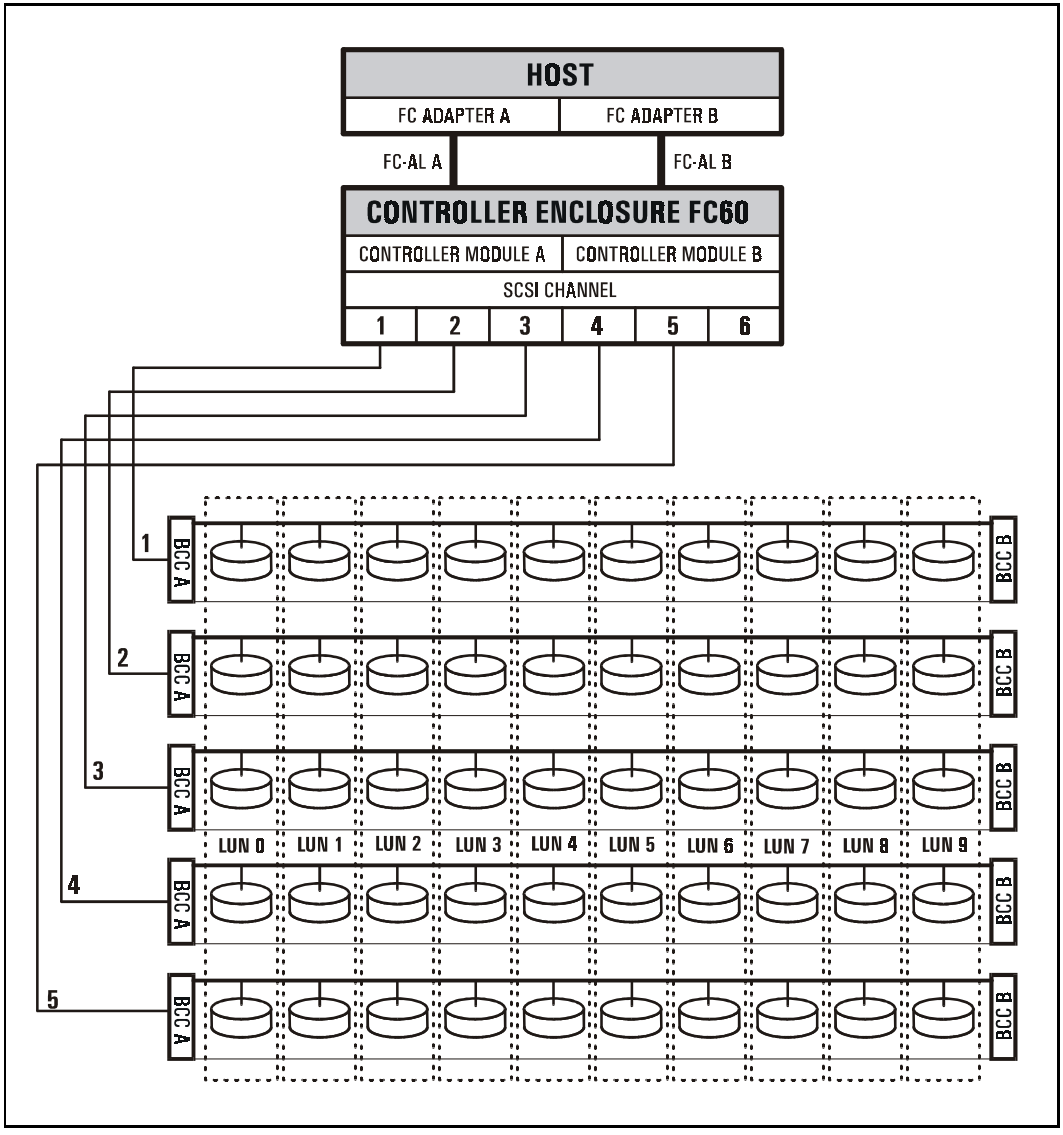



**Figure 32** Five Disk Enclosure Maximum Capacity Configuration

## Six Disk Enclosure Configurations

### High Availability/High Performance

- **Hardware Configuration**
  - Two disk array controllers connected directly to host Fibre Channel adapters
  - Six disk enclosures with ten 73 GByte disks each (60 disks total)
  - Disk enclosures configured for full-bus mode (one SCSI channel per enclosure)
- **LUN Configuration**
  - Ten RAID 0/1 LUNs, each comprising six disks (3+3)
  - Each disk in a LUN is in a separate enclosure
- **High Availability**
  - Handles a single disk failure, single BCC failure, single channel failure, or a single controller failure
- **Disk Capacity**
  - Maximum capacity 4380 GBytes
  - Usable capacity 2190 GBytes

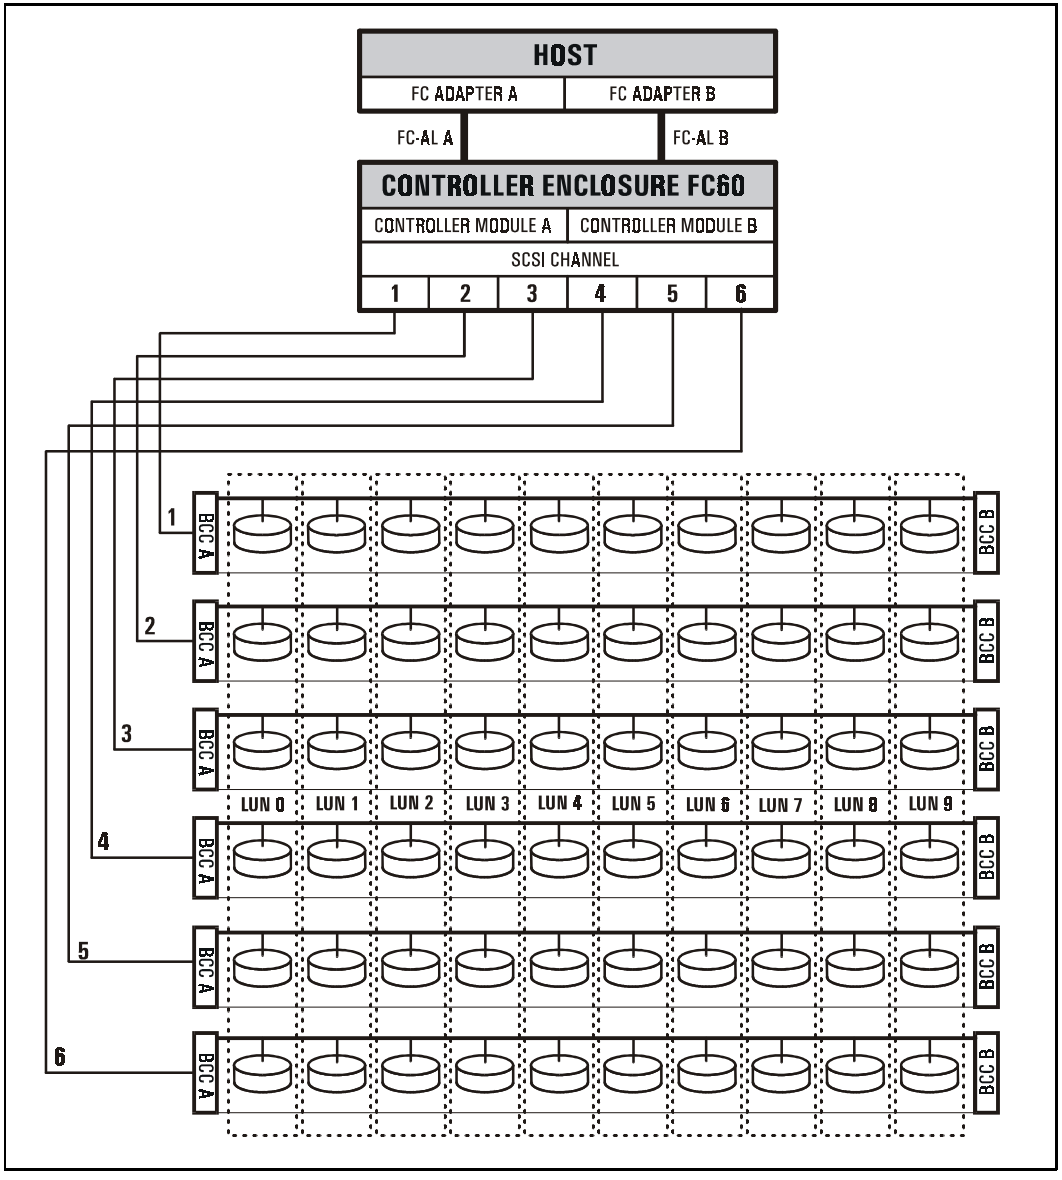



**Figure 33** Six Disk Enclosure High Availability/High Performance Configuration

## Maximum Capacity

- **Hardware Configuration**
  - Two disk array controllers connected directly to host Fibre Channel adapters
  - Six disk enclosures with ten 73 GByte disks each (60 disks total)
  - Disk enclosures configured for full-bus mode (one SCSI channel per enclosure)
- **LUN Configuration**
  - Ten RAID 5 LUNs, each comprising six disks (5 data + 1 parity)
  - Each disk in a LUN is in a separate enclosure
- **High Availability**
  - Handles a single disk failure, single disk enclosure/BCC failure, single channel failure, or a single controller failure
- **Disk Capacity**
  - Maximum capacity 4380 GBytes
  - Usable capacity 3650 GBytes

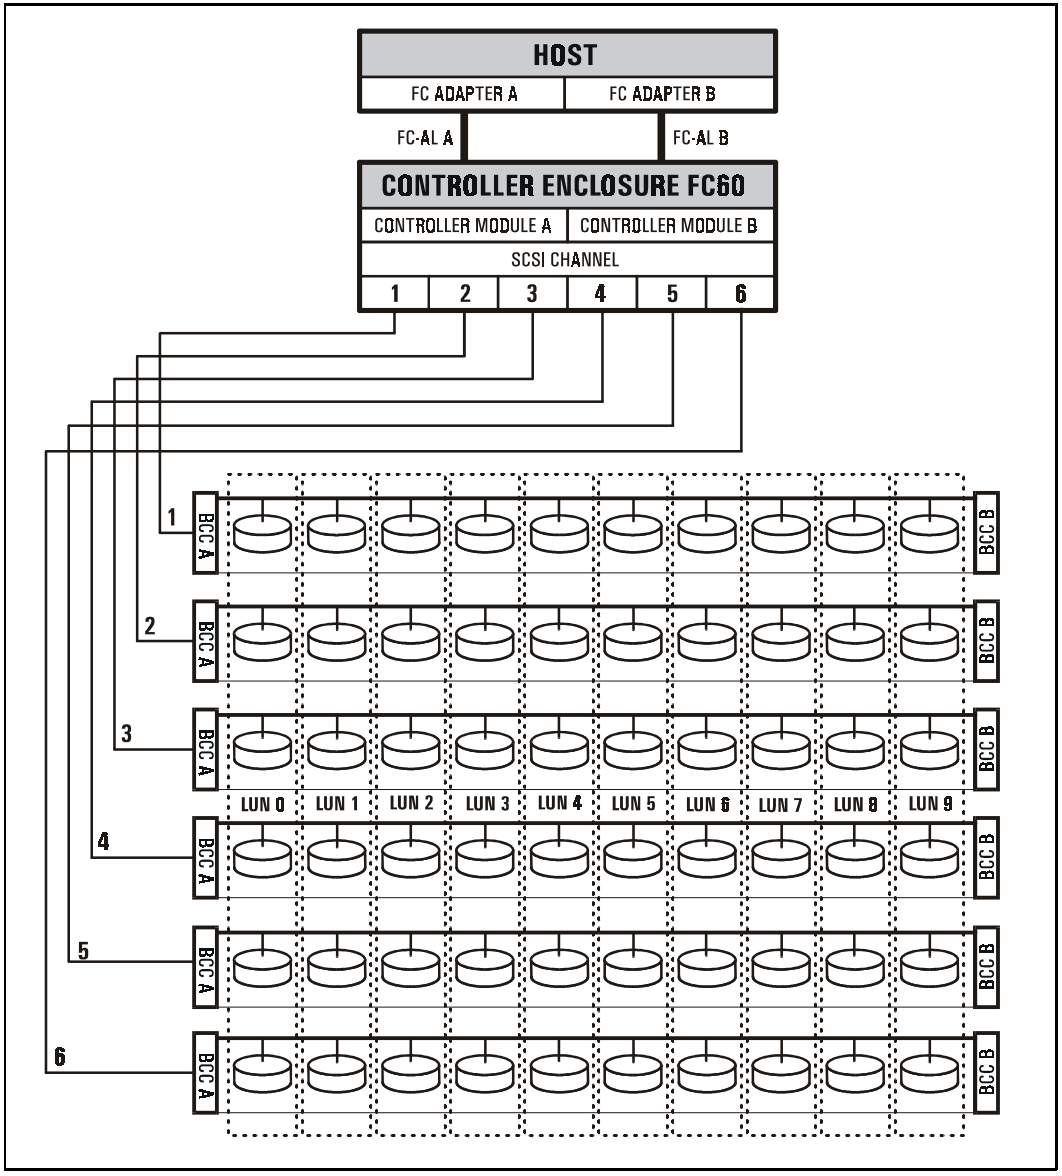



**Figure 34** Six Disk Enclosure High Maximum Capacity Configuration

## Total Disk Array Capacity

The total capacity provided by the disk array depends on the number and capacity of disks installed in the array, and the RAID levels used. RAID levels are selected to optimize performance or capacity.

Table 6 lists the total capacities available when using fully loaded disk enclosures configured for optimum performance. Table 7 lists the same for optimum capacity configurations.

The capacities listed reflect the maximum capacity of the LUN. The actual storage capacity available to the operating system will be slightly less, as some capacity is consumed when binding the LUN and creating the file system.

For high-availability, one disk per SCSI channel is used as a global hot spare.

**Table 6** Capacities for Optimized Performance Configurations

| Number of disk enclosures | RAID Level | No. of LUNs | Disks per LUN | Total Capacity (with indicated disks) | | | |
|---|---|---|---|---|---|---|---|
| | | | | 9.1 GB | 18.2 GB | 36.4 GB | 73 GB |
| 2 (split bus) | 1 | 8 | 2 | 72.8 GB | 145.6 GB | 291.2 GB | 584 GB |
| 3 (split bus) | 1 | 12 | 2 | 109.1 GB | 218.4 GB | 436.8 GB | 876 GB |
| 4 (full bus) | 0/1 | 9 | 4 (2+2) | 163.8 GB | 327.6 GB | 655.2 GB | 1314 GB |
| 6 (full bus) | 0/1 | 9 | 6 (3+3) | 245.7 GB | 491.4 GB | 982.8 GB | 1971 GB |

**Table 7** Capacities for Optimized Capacity Configurations

| Number of disk enclosures | RAID Level | No. of LUNs | Disks per LUN | Total Capacity (with indicated disks) | | | |
|---|---|---|---|---|---|---|---|
| | | | | 9.1 GB | 18.2 GB | 36.2 GB | 73 GB |
| 3 (split bus) | 5 | 8 | 3 (2D+1P) | 145.6 GB | 291.2 GB | 582.4 GB | 1168 GB |
| 4 (full bus) | 5 | 9 | 4 (3D+1P) | 245.7 GB | 491.4 GB | 982.8 GB | 1971 GB |
| 5 (full bus) | 5 | 9 | 5 (4D+1P) | 327.6 GB | 655.2 GB | 1310.4 GB | 2628 GB |
| 6 (full bus) | 5 | 9 | 6 (5D+1P) | 409.5 GB | 819 GB | 1638 GB | 3285 GB |

# Topologies for HP-UX

The topology of a network or a Fibre Channel Arbitrated Loop (Fibre Channel-AL) is the physical layout of the interconnected devices; that is, a map of the connections between the physical devices. The topology of a Fibre Channel-AL is extremely flexible because of the variety and number of devices, or nodes, that can be connected to the Fibre Channel-AL. A node can be a host system or server, a 10-port HP Fibre Channel-AL hub, or the controller modules in a disk array. Two controller modules and two power supply modules are required for high availability topologies. Because it is impossible to document all the possible physical layouts, the hardware topologies supported by the Disk Array FC60 are grouped into five general categories:

- Basic topology
- Single-system distance topology
- High availability topology
- High availability, distance, and capacity topology
- Campus topology

In addition to these topologies which include hubs, the disk array FC60 supports some switch configurations. Each of the five types of topology and switch configurations are documented in the following sections.

## Basic Topology

The basic topology covers a number of physical implementations of host systems and disk arrays. In general, the following hardware components are used:

- One or two host systems
- One Fibre Channel I/O adapter for each implementation of the basic topology in each server, with
- Up to nine adapters in each K -Class server;
  The maximum number depends on the number of HSC slots available in each K -Class server (refer to the Configuration Guide for additional details)
- Up to three adapters installed in any slots including the Turbo EISA slot in each D-Class model D350, D360, D370 and D380 server;
  Up to two adapters installed in any slots in all other models of D-Class server;
- Up to 22 adapters in each T-Class T600 server
- See your Hewlett-Packard support representative for the number of adapters supported in each model of V-Class server
- One or two dual-controller module disk arrays (for high availability) or
  One to four single-controller module disk arrays (for non-high availability)
- Maximum 500 m fibre optic cable distance between each host-to-disk array connection
- Each implementation of the basic topology uses one Fibre Channel I/O adapter. Each adapter connects to one controller module over an independent two-node Fibre Channel-AL. For example, if one disk array with two controller modules connects to one host with two Fibre Channel adapters, two separate two-node Fibre Channel-ALs are created. Up to 22 instances of the basic topology can be implemented on T-Class servers because up to 22 adapters can be installed in each T-Class system.

Supported cable lengths for each host-to-disk array connection include 2 m, 16 m, 50 m, 100 m, and 500 m. Fibre optic cables in lengths of 2 m, 16 m, 50 m, and 100 m can be ordered from Hewlett-Packard (see the chapter 8 for part numbers). Fibre optic cables longer than 100 m must be custom-fabricated for each implementation.

Non-high availability and high availability versions of this topology can be implemented depending on whether one or two controller modules are installed in each disk array.

For high availability the hosts and disk arrays can be connected in any of the following ways, with each connection of adapter and controller module creating a separate 2-node Fibre Channel-AL:

- One disk array with two controller modules with each controller module connected to a separate adapter in a single host
- In K-Class and T-Class systems, two disk arrays with two controller modules per array, connected to a single host with four Fibre Channel adapters.

The high availability version uses disk arrays having two (redundant) controller modules. Each controller module is connected to a separate adapter. The two adapters can be installed in one host or each adapter can be installed in a different host. Disk arrays having redundant controller modules protect against single points of hardware failure by providing an alternate hardware path to the host if an controller module, Fibre Channel cable, host Fibre Channel I/O adapter, or internal Ultra2 SCSI bus fails. Applications can continue to run after the failure of a single disk module. This is explained in the “Disk Array High Availability Features” section in chapter 1.

Figure 35 shows the high availability version of the basic topology implemented on a host system with two host Fibre Channel I/O adapters connected to a dual-controller module disk array. One high availability version of the basic topology can be implemented in D-Class servers.

**Note** Connecting multiple hosts to the Disk Array FC60 requires the use of Fibre Channel-AL hubs. Connecting the host through hubs preserves high-availability in the event of a host adapter failure.

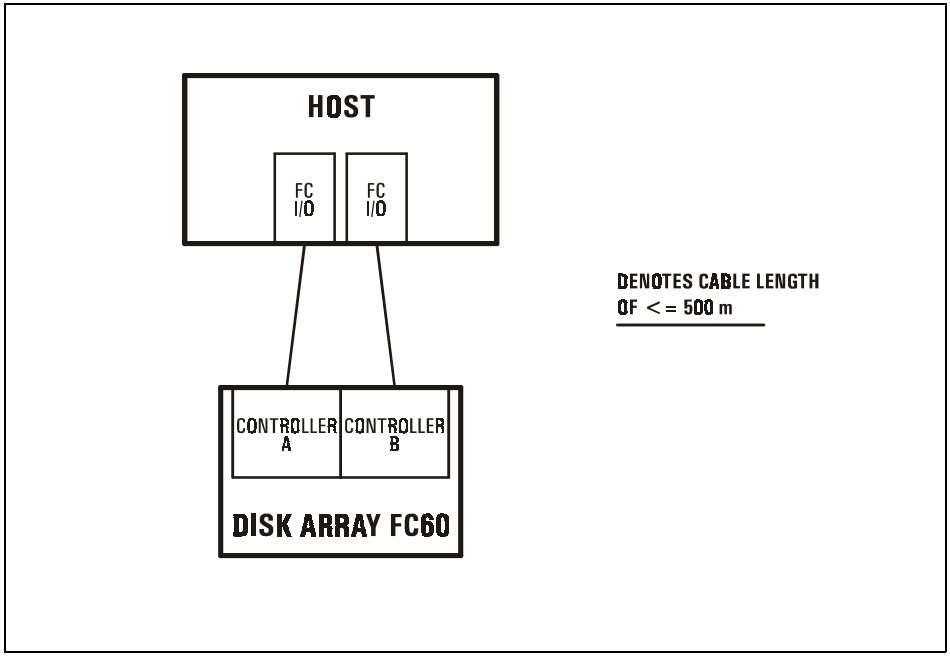



**Figure 35** Basic Topology, High Availability Version: Host with Two Fibre Channel I/O Adapters

Figure 36 shows the high availability version of the basic topology implemented on a either a K -Class or T-Class host with four Fibre Channel I/O adapters. Two of the Fibre Channel adapters are connected to one dual-controller module disk array while the other two Fibre Channel adapters are connected to a second dual-controller module disk array. Each connection of host adapter and controller module creates a separate Fibre Channel-AL.

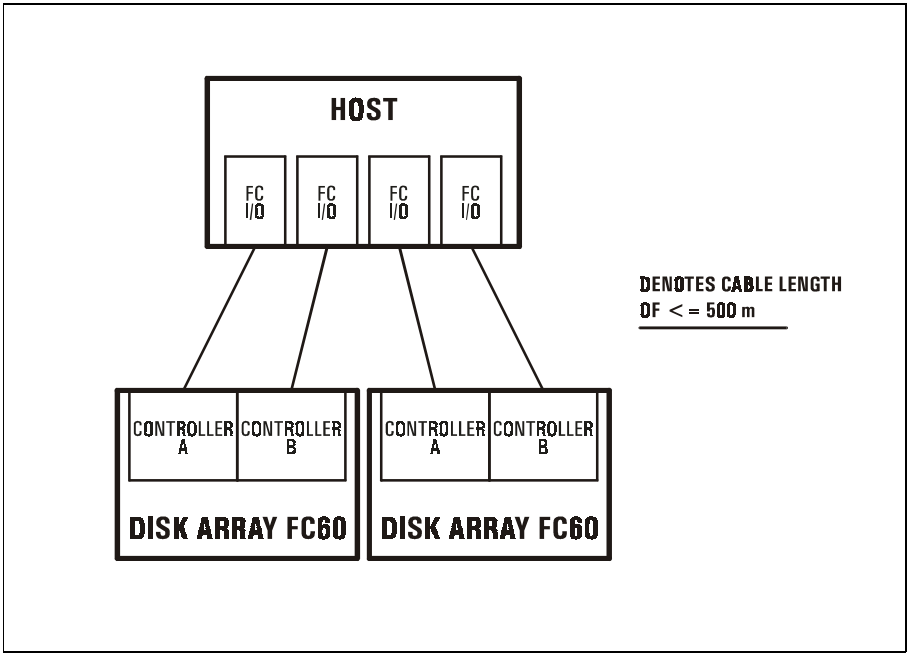



**Figure 36** Basic Topology, High Availability Version: Host with Four Fibre Channel I/O Adapters

The non-high availability version of this topology connects a host or server to one or more single-controller module disk arrays. This version provides no hardware redundancy and does not protect against single points of controller module, Fibre Channel cable, host Fibre Channel I/O adapter, or internal Ultra2 SCSI bus failure. If any of these components fail, the disk array becomes unavailable and applications cannot continue to run. The disk array remains unavailable until the failed hardware component is replaced. Applications can continue to run after the failure of single disk modules within logical disk units (LUNs).

Figure 37 shows the non-high availability version of the basic topology implemented on either a K -Class or T-Class host with four Fibre Channel I/O adapters connected to single

controller modules in four disk arrays. Each connection between adapter and controller module creates a separate Fibre Channel-AL.

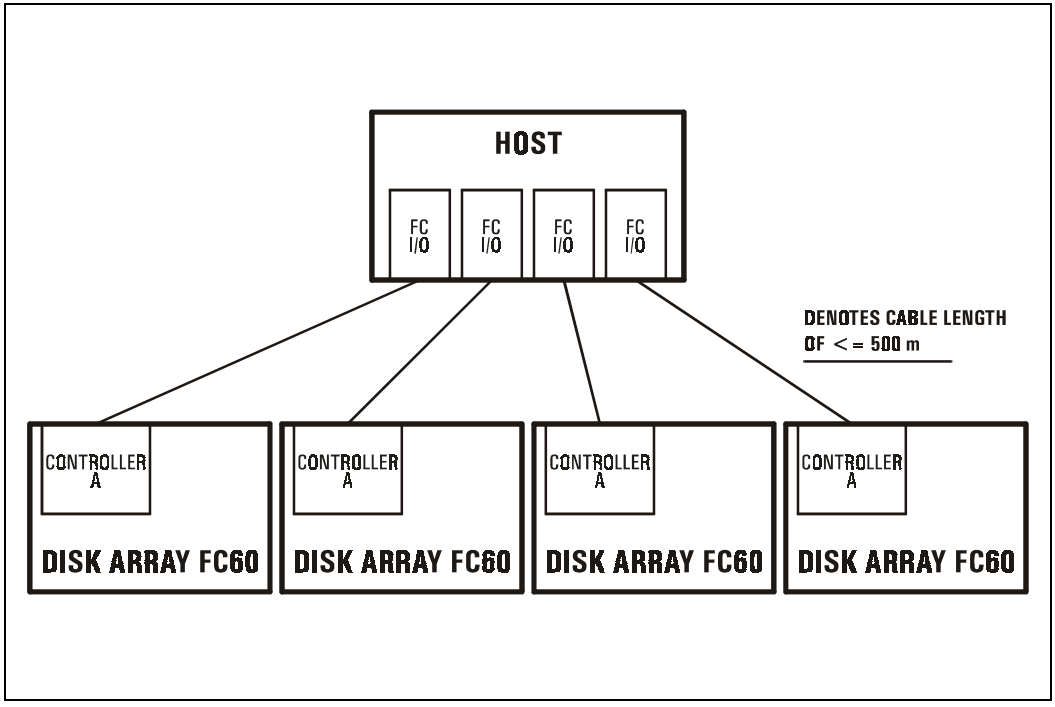



**Figure 37** Basic Topology, Non-High Availability Version: Host with Four Fibre Channel I/O Adapters

**Table 8** Basic Topology Error Recovery

| Failing component | Continue after failure | What happens and how to recover |
|---|---|---|
| Disk module | Yes | Applications continue to run on all supported RAID levels (RAID-1, 0/1, and 5). The system administrator or service provider hot-swaps the failed disk module. |
| Controller module in single-controller module array (no high availability) | No | The disk array fails. The system administrator or service provider replaces the controller module and the system administrator restarts the operating system and applications. |
| Controller module in dual controller module disk array | No on path to failed controller module;<br>Yes on alternate path through second controller module | Ownership of the failed controller module's LUNs can transfer automatically to the remaining operational controller module if primary and alternate paths have been configured in LVM. If so, LVM switches automatically and transparently to the alternate path. The system administrator or service provider hot-swaps the controller module. |
| Fan module | Yes | Applications continue to run. The system administrator or service provider hot-swaps the fan module. |
| Power supply modules | Yes | Applications continue to run when one power supply module fails. The system administrator or service provider hot-swaps the failed power supply module. |
| Fibre Channel I/O adapter | No on path to failed adapter;<br>Yes if array has dual controller modules and alternate paths have been configured | I/O operations fail along the path through the failed adapter. If the host has two Fibre Channel adapters connected to a dual controller module disk array, the array can be accessed through the path to the operational adapter and controller module if primary and alternate paths have been configured in LVM. If so, LVM switches automatically and transparently to the alternate path. The authorized service provider replaces the failed adapter and the system administrator restarts the operating system and applications. |

**Table 8** Basic Topology Error Recovery (cont'd)

| Failing component | Continue after failure | What happens and how to recover |
|---|---|---|
| Fibre Channel cable | No on path to failed cable; Yes if array has dual controller modules and alternate paths have been configured | I/O operations fail along the path through the failed cable. If the host has two Fibre Channel adapters connected to a dual-controller module disk array, the array can be accessed through the operational path if primary and alternate paths have been configured in LVM. If so, LVM switches automatically and transparently to the alternate path. The authorized service provider hot-swaps the failed cable. |

## Single-System Distance Topology

Each instance of the single-system distance topology generally uses the following hardware components:

- One host system or server
- Two Fibre Channel I/O adapters in each server
- One 10-port HP Fibre Channel-AL Hub
- One to four dual-controller module disk arrays (for high availability) or
  One to eight single-controller module disk arrays (for non-high availability)
- Maximum 500 m fibre optic cable distance on each connection between the host and the HP Fibre Channel-AL Hub and between the HP Fibre Channel-AL Hub and each disk array, for a total Fibre Channel-AL cable length not to exceed 5000 m

One single-system distance topology can be implemented in each D-Class system, one or two single-system topologies can be implemented in each K-Class system, and up to 11 instances of this topology can be implemented in each T-Class server. See your Hewlett-Packard support representative for the number of adapters and topologies that can be implemented in V-Class systems.

This topology uses one 10-port HP Fibre Channel-AL Hub. The two Fibre Channel I/O adapters in the single host attach to two ports in the HP Fibre Channel-AL Hub, providing redundant hardware paths between the host and the HP Fibre Channel-AL Hub. The controller modules of dual or single-controller module disk arrays attach to some or all of the remaining eight HP Fibre Channel-AL Hub ports. This topology creates a single Fibre Channel-AL that resembles a star topology with the HP Fibre Channel-AL Hub acting as the central switching element. The 10-port HP Fibre Channel-AL Hub is inserted between the host and the disk arrays, increasing both the number of arrays (and total disk capacity) that can be connected to a single host and the total length of the Fibre Channel cabling. The HP Fibre Channel-AL Hub is a single point of failure in this topology.

Supported cable lengths for each segment of the Fibre Channel-AL include 2 m, 16 m, 50 m, 100 m, and 500 m. The maximum combined cable lengths for all segments, that is, the total length of the Fibre Channel-AL should not exceed 5000 m because performance can degrade due to propagation delay. Because of this it is recommended that the total cable length of the Fibre Channel-AL be as short as possible. Fibre optic cables in lengths of 2 m, 16 m, 50 m, and 100 m cables can be ordered from Hewlett-Packard (see the Reference

chapter for part numbers). Fibre optic cables longer than 100 m must be custom-fabricated for each implementation.

Like the basic topology, both high availability versions (two controller modules per disk array) and non-high availability versions (one controller module per disk array) of this topology can be implemented.

For high availability implementations, up to four disk arrays with two controller modules per disk array can be connected to an HP Fibre Channel-AL Hub. The two Fibre Channel I/O adapters installed in the host connect to the two remaining HP Fibre Channel-AL Hub ports. If one controller module fails, ownership of its LUNs can transfer automatically to the remaining operational controller module through the other Fibre Channel adapter. If a primary and an alternate path have been configured in LVM, LVM can switch automatically and transparently to the alternate path. Likewise, if one host adapter fails, or if one Fibre Channel cable from the host to the HP Fibre Channel-AL Hub or from the HP Fibre Channel-AL Hub to one of the dual controller module disk array fails, that disk array can still be accessed via the alternate hardware path. The HP Fibre Channel-AL Hub is a single point of failure in this topology. If the HP Fibre Channel-AL Hub fails, no communication between the host and any of the disk arrays is possible.

For non-high availability, up to eight single-controller module disk arrays can be attached to the HP Fibre Channel-AL Hub. If the single controller module or the Fibre Channel cable between the HP Fibre Channel-AL Hub and an array fails, no I/O operations are possible between the host and the disk array with the failed controller module. If one Fibre Channel adapter fails in the host, the disk array can still be accessed via the hardware path through the other operational Fibre Channel adapter. If the HP Fibre Channel-AL Hub fails, no communication between the host and any disk arrays is possible.

Figure 38 illustrates the single-system distance topology with one host with two Fibre Channel I/O adapters and three dual-controller module disk arrays. In this example two of the HP Fibre Channel-AL Hub's ten ports are unused.

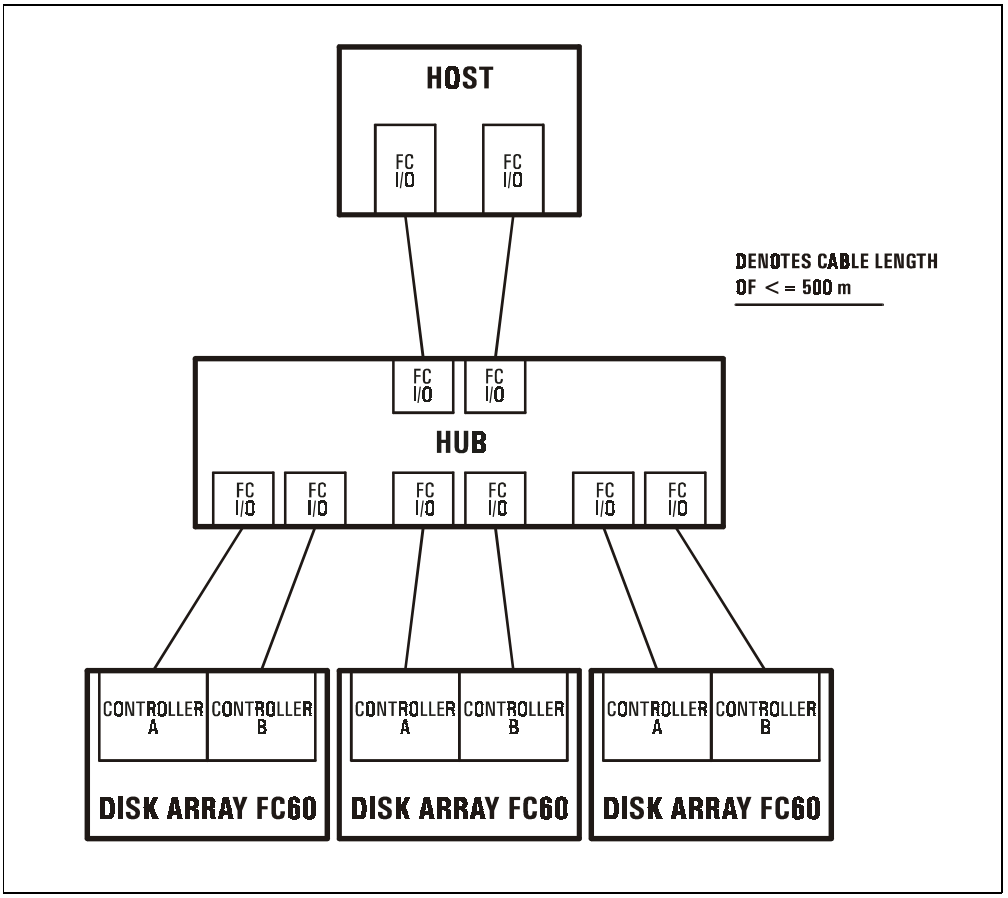



**Figure 38** Single-System Distance Topology

**Table 9** Single-System Distance Topology Error Recovery

| Failing component | Continue after failure | What happens and how to recover |
|---|---|---|
| Disk module | Yes | Applications continue to run on all supported RAID levels (RAID-1, 0/1, and 5). The system administrator or service provider hot-swaps the failed disk module. |
| Controller module in single-controller module array (no high availability) | No | The disk array fails. The system administrator or service provider replaces the controller module and the system administrator restarts the operating system and applications. |
| Controller module in dual-controller module disk array | No on path to failed controller module; Yes on alternate path through second controller module | Ownership of the failed controller module's LUNs can transfer automatically to the remaining operational controller module if primary and alternate paths have been configured in LVM. If so, LVM switches automatically and transparently to the alternate path. The system administrator or service provider hot-swaps the failed controller module. |
| Fan module | Yes | Applications continue to run. The system administrator or service provider hot-swaps the fan module. |
| Power supply modules | Yes | Applications continue to run when one power supply module fails. The system administrator or service provider hot-swaps the failed power supply module. |
| HP Fibre Channel-AL Hub | No | No I/O operations are possible between the host and any disk arrays because all I/O traffic is routed through the HP Fibre Channel-AL Hub. The authorized service provider replaces the failed HP Fibre Channel-AL Hub and the system administrator restarts the operating system and applications. |

**Table 9** Single-System Distance Topology Error Recovery (cont'd)

| Failing component | Continue after failure | What happens and how to recover |
|---|---|---|
| Fibre Channel I/O adapter | No on path to failed adapter; Yes if array has dual controller modules and alternate paths have been configured | I/O operations fail along the path through the failed adapter. If the host has two Fibre Channel adapters connected to a dual-controller module disk array, the array can be accessed through the path to the operational adapter and controller module if primary and alternate paths have been configured in LVM. If so, LVM switches automatically and transparently to the alternate path. The authorized service provider replaces the failed adapter and the system administrator restarts the operating system and applications. |
| Fibre Channel cable | No on path to failed cable; Yes if array has dual controller modules and alternate paths have been configured | I/O operations fail along the path through the failed cable. If the host has two Fibre Channel adapters connected to a dual-controller module disk array, the array can be accessed through the operational path if primary and alternate paths have been configured in LVM. If so, LVM switches automatically and transparently to the alternate path. The authorized service provider hot-swaps the failed cable. |

## High Availability Topology

The high availability topology increases the availability of the single system distance topology by protecting against single points of HP Fibre Channel-AL Hub failure with the use of redundant HP Fibre Channel-AL Hubs connected to each other. Adding a second HP Fibre Channel-AL Hub also increases the number of hosts and disk arrays that can be connected to a single Fibre Channel-AL.

Each instance of the high availability topology uses the following hardware components:

- Two to four host systems or servers
- Two Fibre Channel I/O adapters per host
- Two 10-port HP Fibre Channel-AL Hubs
- Two controller modules per disk array
- Maximum of six dual-controller module disk arrays in a four-host Fibre Channel-AL or Maximum of eight dual-controller module disk arrays in a two-host Fibre Channel-AL
- Maximum 500 m fibre optic cable distance on each connection between each host and HP Fibre Channel-AL Hub and between each HP Fibre Channel-AL Hub and disk array, for a total Fibre Channel-AL cable length not to exceed 5000 m

As its name implies, the high availability topology supports only dual-controller module disk arrays (Figure 39). The objective of this topology is to use redundant hardware components to provide high availability, so disk arrays with single controller modules are not supported. All hardware components and paths are redundant. If any hardware failure occurs, I/O communication between host and disk array can be completed through another path.

In this topology each host uses two Fibre Channel I/O adapters and each disk array uses dual controller modules. One adapter in each host and one controller module in each disk array connect to one HP Fibre Channel-AL Hub, forming one Fibre Channel-AL. The second adapter in each host and the second controller module in each disk array connect to the second HP Fibre Channel-AL Hub, cabled into a separate Fibre Channel-AL. This creates two redundant Fibre Channel-ALs.

Because each HP Fibre Channel-AL Hub has ten ports, either two host adapters and eight controller modules or four host adapters and six controller modules can attach to each HP Fibre Channel-AL Hub. If any hardware component (controller module, host Fibre Channel I/O adapter, HP Fibre Channel-AL Hub, or fibre optic cable) fails in one Fibre Channel-AL, the I/O communication between hosts and disk arrays can continue through the other Fibre Channel-AL.

Supported cable lengths for each segment of the Fibre Channel-AL include 2 m, 16 m, 50 m, 100 m, and 500 m. The maximum combined cable lengths for all segments, that is, the total length of the Fibre Channel-AL should not exceed 5000 m because performance can degrade due to propagation delay. Because of this it is recommended that the total cable length of the Fibre Channel-AL be as short as possible. Fibre optic cables in lengths of 2 m, 16 m, 50 m, and 100 m cables can be ordered from Hewlett-Packard (see chapter 8 for part numbers). Fibre optic cables longer than 100 m must be custom-fabricated for each implementation.

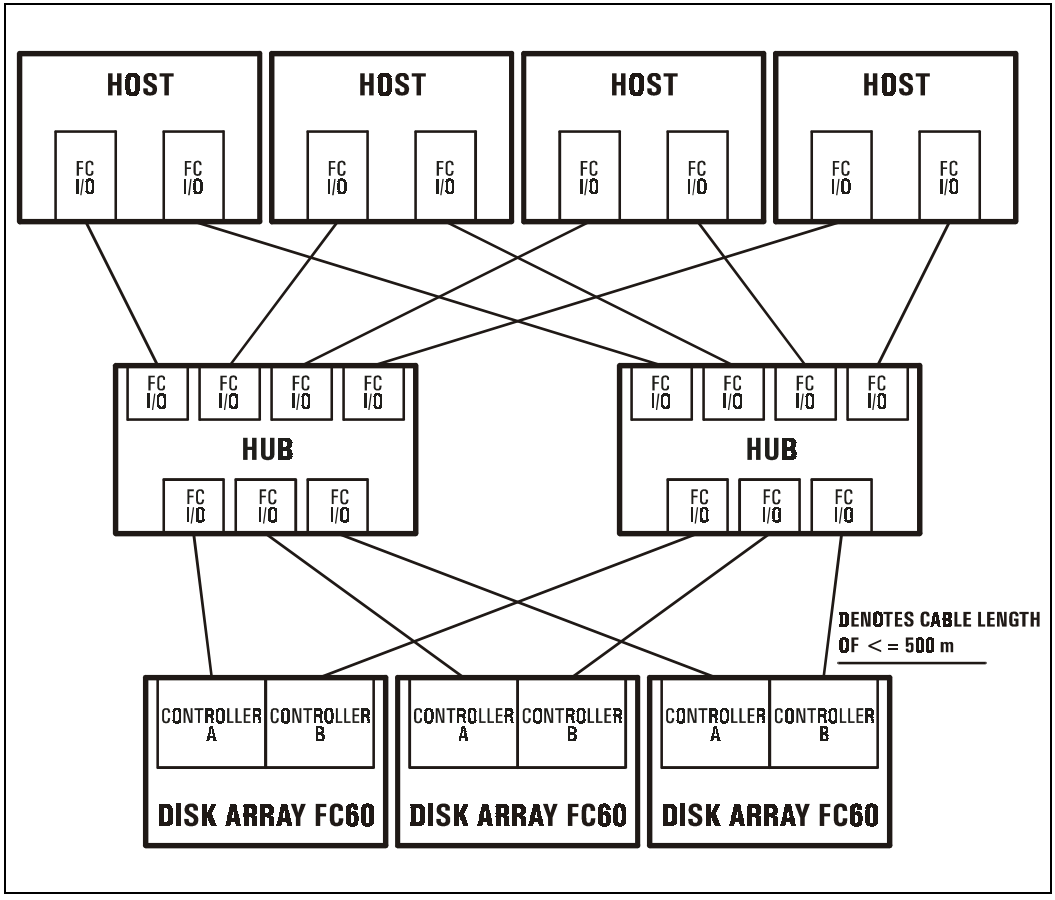



**Figure 39** High Availability Topology

**Table 10** High Availability Topology Error Recovery

| Failing component | Continue after failure | What happens and how to recover |
|---|---|---|
| Disk module | Yes | Applications continue to run on all supported RAID levels (RAID-1, 0/1, and 5). The system administrator or service provider hot-swaps the failed disk module. |
| Controller modules | No on path to failed controller module; Yes on alternate path through second controller module | Ownership of the failed controller module's LUNs can transfer automatically to the remaining operational controller module if primary and alternate paths have been configured in LVM. If so, LVM switches automatically and transparently to the alternate path. The system administrator or service provider hot-swaps the failed controller module. |
| Fan module | Yes | Applications continue to run. The system administrator or service provider hot-swaps the fan module. |
| Power supply modules | Yes | Availability are present, applications continue to run when one power supply module fails. The system administrator or service provider hot-swaps the failed power supply module. |
| HP Fibre Channel-AL Hub | No on path to failed hub; Yes on alternate path to other hub | I/O operations fail along the path to the failed HP Fibre Channel-AL Hub. I/O operations can transfer automatically to the other HP Fibre Channel-AL Hub if primary and alternate paths have been configured in LVM. If so, LVM switches automatically and transparently to the alternate path. The authorized service provider hot-swaps the failed HP Fibre Channel-AL Hub. |
| Fibre Channel I/O adapter | No on path to failed adapter; Yes if array has dual controller modules and alternate paths have been configured | I/O operations fail along the path through the failed adapter. If the host has two Fibre Channel adapters connected to a dual-controller module disk array, the array can be accessed through the path to the operational adapter and controller module if primary and alternate paths have been configured in LVM. If so, LVM switches automatically and transparently to the alternate path. The authorized service provider replaces the failed adapter and the system administrator restarts the operating system and applications. |

**Table 10** High Availability Topology Error Recovery (cont'd)

| Failing component | Continue after failure | What happens and how to recover |
|---|---|---|
| Fibre Channel cable | No on path to failed cable; Yes if array has dual controller modules and alternate paths have been configured | I/O operations fail along the path through the failed cable. If the host has two Fibre Channel adapters connected to a dual-controller module disk array, the array can be accessed through the operational path if primary and alternate paths have been configured in LVM. If so, LVM switches automatically and transparently to the alternate path. The authorized service provider hot-swaps the failed cable. |

## High Availability, Distance, and Capacity Topology

The high availability, distance, and capacity topology expands on the high availability topology by using cascaded HP Fibre Channel-AL Hubs to increase the distance of each Fibre Channel-AL and the number of devices that can be interconnected on the Fibre Channel-AL. Cascaded HP Fibre Channel-AL Hubs are two HP Fibre Channel-AL Hubs connected together.

Each instance of this topology uses the following hardware components:

- Up to four host systems or servers
- Two Fibre Channel I/O adapters per host
- Two pairs of cascaded 10-port HP Fibre Channel-AL Hubs
- Two controller modules per disk array
- Maximum of nine dual-controller module disk arrays
- Maximum 500 m fibre optic cable distance on each connection between each host and HP Fibre Channel-AL Hub and between each HP Fibre Channel-AL Hub and disk array, for a total Fibre Channel-AL cable length not to exceed 5000 m

Like the high availability topology, this topology supports high availability by using redundant Fibre Channel-ALs. If a hardware component (Fibre Channel adapter, HP Fibre Channel-AL Hub, controller module or cables) in one Fibre Channel-AL fails, I/O communication between hosts and disk arrays can continue through the other Fibre Channel-AL.

The increased distance is supported by using 500 m fibre optic cable to connect each pair of cascaded HP Fibre Channel-AL Hubs. If distance is a requirement, it is managed between the two HP Fibre Channel-AL Hubs. The distance from hosts to HP Fibre Channel-AL Hubs and from disk arrays to HP Fibre Channel-AL Hubs should be minimized.

Supported cable lengths for each segment of the Fibre Channel-AL include 2 m, 16 m, 50 m, 100 m, and 500 m. The maximum combined cable lengths for all segments, that is, the total length of the Fibre Channel-AL should not exceed 5000 m because performance can degrade due to propagation delay. Because of this it is recommended that the total cable length of the Fibre Channel-AL be as short as possible. Fibre optic cables in lengths of 2 m, 16 m, 50 m, and 100 m cables can be ordered from Hewlett-Packard (see the Reference chapter for part numbers). Fibre optic cables longer than 100 m must be custom-fabricated for each implementation.

In this configuration a maximum of nine controller modules can be attached to an HP Fibre Channel-AL Hub. The HP Fibre Channel-AL Hub's tenth port is used to connect to the other HP Fibre Channel-AL Hub in the cascaded pair, rather than to a host adapter as in the high availability topology.

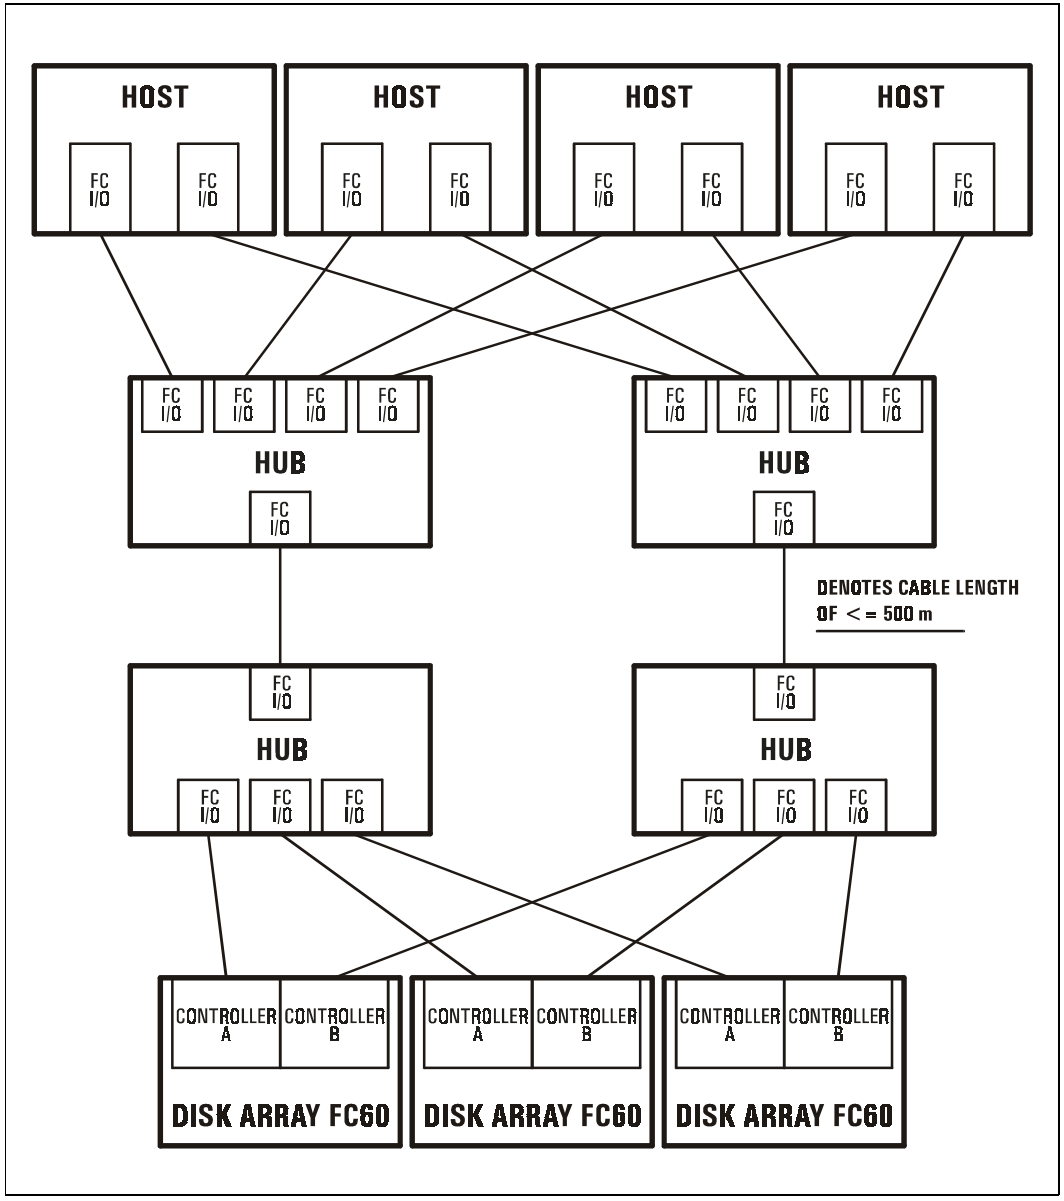



**Figure 40** High Availability, Distance, and Capacity Topology

**Table 11** High Availability, Distance, and Capacity Topology Error Recovery

| Failing component | Continue after failure | What happens and how to recover |
|---|---|---|
| Disk module | Yes | Applications continue to run on all supported RAID levels (RAID-1, 0/1, and 5). The system administrator or service provider hot-swaps the failed disk module. |
| Controller modules | No on path to failed controller module; Yes on alternate path through second controller module | Ownership of the failed controller module's LUNs can transfer automatically to the remaining operational controller module if primary and alternate paths have been configured in LVM. If so, LVM switches automatically and transparently to the alternate path. The system administrator or service provider hot-swaps the failed controller module. |
| Fan module | Yes | Applications continue to run. The system administrator or service provider hot-swaps the fan module. |
| Power supply modules | Yes | Applications continue to run when one power supply module fails. The system administrator or service provider hot-swaps the failed power supply module. |
| HP Fibre Channel-AL Hub | No on path to failed hub; Yes on alternate path to other pair of hubs | I/O operations fail along the path to the failed HP Fibre Channel-AL Hub. I/O operations can transfer automatically to the other pair of cascaded HP Fibre Channel-AL Hubs if primary and alternate paths have been configured in LVM. If so, LVM switches automatically and transparently to the alternate path. The authorized service provider hot-swaps the failed HP Fibre Channel-AL Hub. |
| Fibre Channel I/O adapter | No on path to failed adapter; Yes if array has dual controller modules and alternate paths have been configured | I/O operations fail along the path through the failed adapter. If the host has two Fibre Channel adapters connected to a dual-controller module disk array, the array can be accessed through the path to the operational adapter and controller module if primary and alternate paths have been configured in LVM. If so, LVM switches automatically and transparently to the alternate path. The authorized service provider replaces the failed adapter and the system administrator restarts the operating system and applications. |

**Table 11** High Availability, Distance, and Capacity Topology Error Recovery (cont'd)

| Failing component | Continue after failure | What happens and how to recover |
|---|---|---|
| Fibre Channel cable | No on path to failed cable; Yes if array has dual controller modules and alternate paths have been configured | I/O operations fail along the path through the failed cable. If the host has two Fibre Channel adapters connected to a dual-controller module disk array, the array can be accessed through the operational path if primary and alternate paths have been configured in LVM. If so, LVM switches automatically and transparently to the alternate path. The authorized service provider hot-swaps the failed cable. |

## Campus Topology

The campus topology uses the same hardware components as the high availability, distance, and capacity topology. The components for each instance of this topology include:

- Up to four host systems or servers
- Two Fibre Channel I/O adapters per host
- Two pairs of cascaded 10-port HP Fibre Channel-AL Hubs
- Two controller modules per disk array
- Maximum of nine dual-controller module disk arrays
- Maximum 500 m fibre optic cable distance on each connection between each host and HP Fibre Channel-AL Hub and between each HP Fibre Channel-AL Hub and disk array, for a total Fibre Channel-AL cable length not to exceed 5000 m

This topology is almost identical to the high availability, distance, and capacity topology. The difference is that in the campus topology one-half of the host systems, HP Fibre Channel-AL Hubs, and disk arrays are located in one building while the other half of the hosts, HP Fibre Channel-AL Hubs, and disk arrays are located at another site. The distribution of systems and hardware between two physical sites provides protection against a power failure or some other catastrophic site-wide failure.

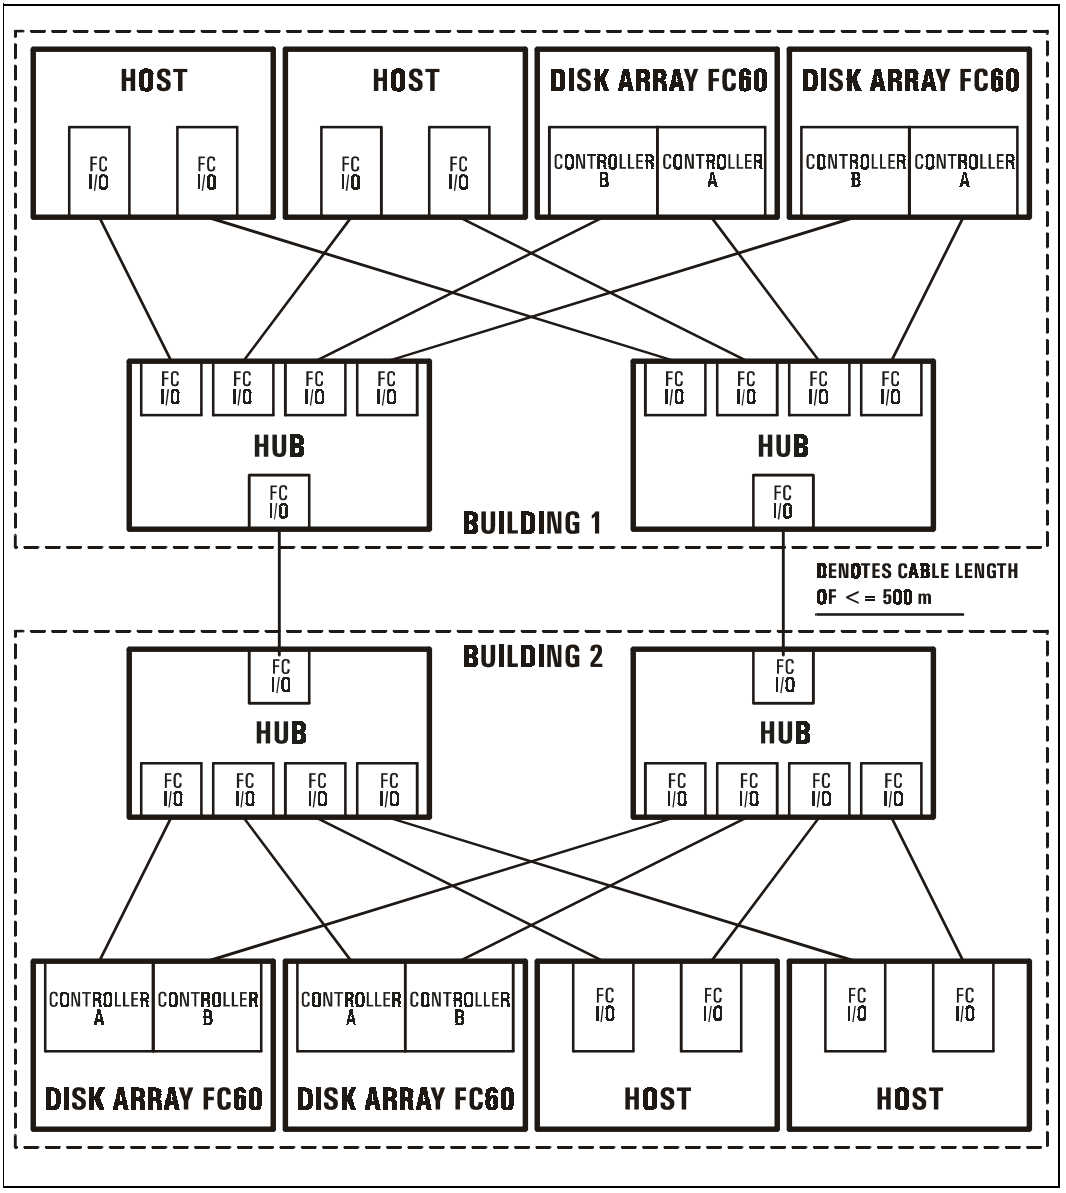



**Figure 41** Campus Topology

**Table 12** Campus Topology Error Recovery

| Failing component | Continue after failure | What happens and how to recover |
|---|---|---|
| Disk module | Yes | Applications continue to run on all supported RAID levels (RAID-1, 0/1, and 5). The system administrator or service provider hot-swaps the failed disk module. |
| Controller module | No on path to failed controller module; Yes on alternate path through second controller module | Ownership of the failed controller module's LUNs can transfer automatically to the remaining operational controller module if primary and alternate paths have been configured in LVM. If so, LVM switches automatically and transparently to the alternate path. The system administrator or service provider hot-swaps the failed controller module. |
| Fan module | Yes | Applications continue to run. The system administrator or service provider hot-swaps the fan module. |
| Power supply modules | Yes | Applications continue to run when one power supply modules fails. The system administrator or service provider hot-swaps the failed power supply module. |
| HP Fibre Channel-AL Hub | No on path to failed hub; Yes on alternate path to other pair of hubs | I/O operations fail along the path to the failed HP Fibre Channel-AL Hub. I/O operations can transfer automatically to the other pair of cascaded HP Fibre Channel-AL Hubs if primary and alternate paths have been configured in LVM. If so, LVM switches automatically and transparently to the alternate path. The authorized service provider hot-swaps the failed HP Fibre Channel-AL Hub. |
| Fibre Channel I/O adapter | No on path to failed adapter; Yes if array has dual controller modules and alternate paths have been configured | I/O operations fail along the path through the failed adapter. If the host has two Fibre Channel adapters connected to a dual-controller module disk array, the array can be accessed through the path to the operational adapter and controller module if primary and alternate paths have been configured in LVM. If so, LVM switches automatically and transparently to the alternate path. The authorized service provider replaces the failed adapter and the system administrator restarts the operating system and applications. |

| Failing component | Continue after failure | What happens and how to recover |
|---|---|---|
| Fibre Channel cable | No on path to failed cable; Yes if array has dual controller modules and alternate paths have been configured | I/O operations fail along the path through the failed cable. If the host has two Fibre Channel adapters connected to a dual-controller module disk array, the array can be accessed through the operational path if primary and alternate paths have been configured in LVM. If so, LVM switches automatically and transparently to the alternate path. The authorized service provider hot-swaps the failed cable |

## Performance Topology with Switches

Previous topologies use Fibre Channel HUBs for interconnecting the arrays with the hosts. In these topologies there is basically one loop with all devices connected to it. The disk array FC60 can be connected to switches. Connecting the array using switches provides for increased performance. Two switch topologies are provided in Figure 42 and Figure 43.

**Note** Hubs and switches should not be mixed in these array topologies. All devices must be directly connected to a switch.

Switches require firmware version V.6C2 to support the disk array FC60

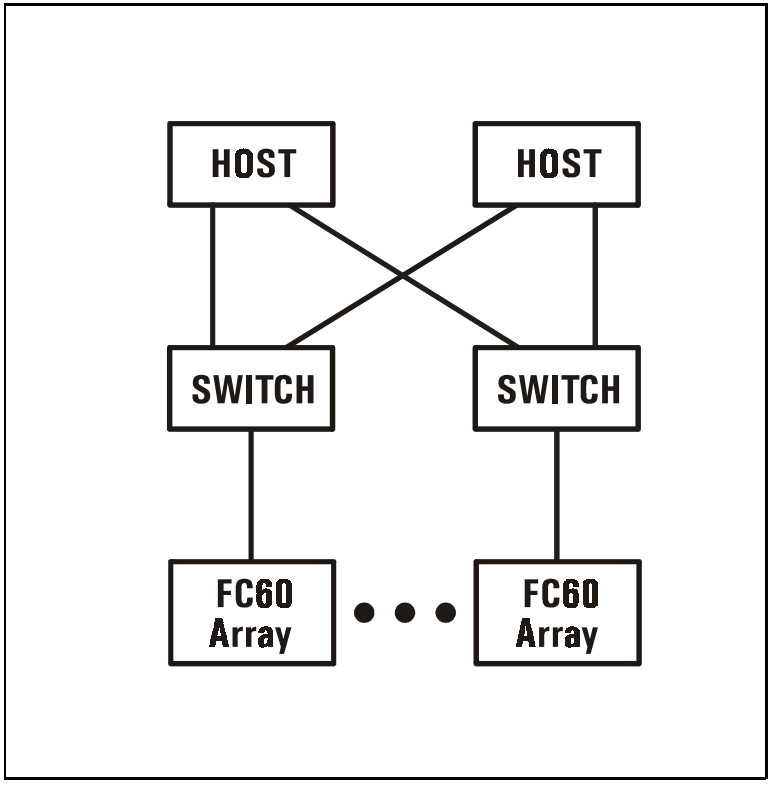



**Figure 42** Two Hosts Connected to Redundant Switches

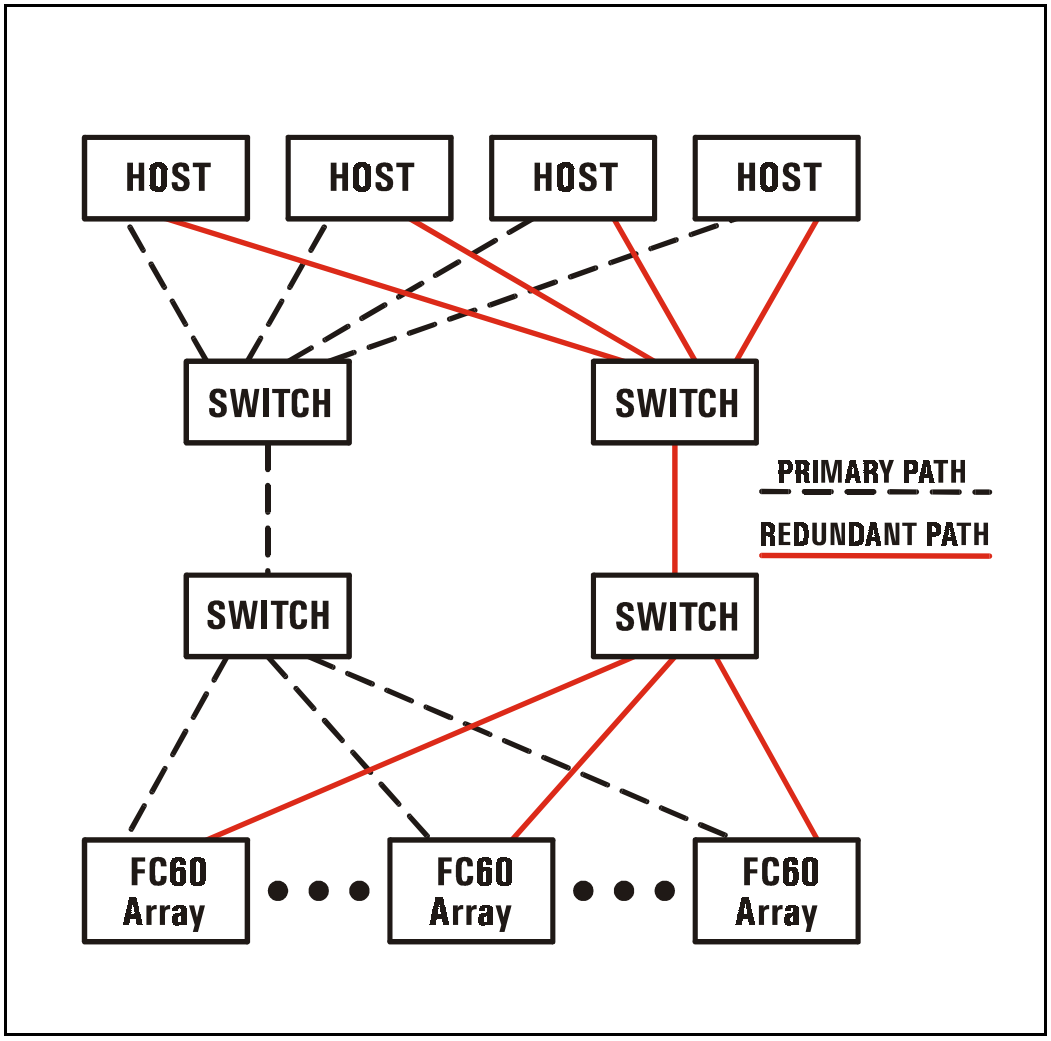



**Figure 43** Four Hosts Connected to Cascaded Switches

# Topologies for Windows NT and Windows 2000

The topology of a network or a Fibre Channel Arbitrated Loop (Fibre Channel-AL) is the physical layout of the interconnected devices; that is, a map of the connections between the physical devices. The topology of a Fibre Channel-AL is extremely flexible because of the variety and number of devices, or nodes, that can be connected to the Fibre Channel-AL. A node can be a host system or server, a Fibre Channel hub or switch, or the controller modules in a disk array.

The tested topologies supported on Windows NT and Windows 2000 are illustrated in this section. Included are the error handling capabilities of the topologies.

**Note**

**Unsupported Single Hub Topologies on Windows.** Due to limitations in the RDAC driver included with the *HP Storage Manager 60* (SM60) software, some topologies that use a single hub are not supported. The RDAC driver supports only two data paths to the FC60 Disk Array, so any topology that provides more than two paths will not work. This is most common is topologies using a single hub. Topologies using a Fibre Channel switch are supported.

Any topology that is configured in the following manner is not supported with the SM60 software:

**A host with dual Fibre Channel adapters both connected to a single hub, with the hub then connected to both controllers on the disk array.**

Figure 44 shows a sample of an unsupported Windows topology.

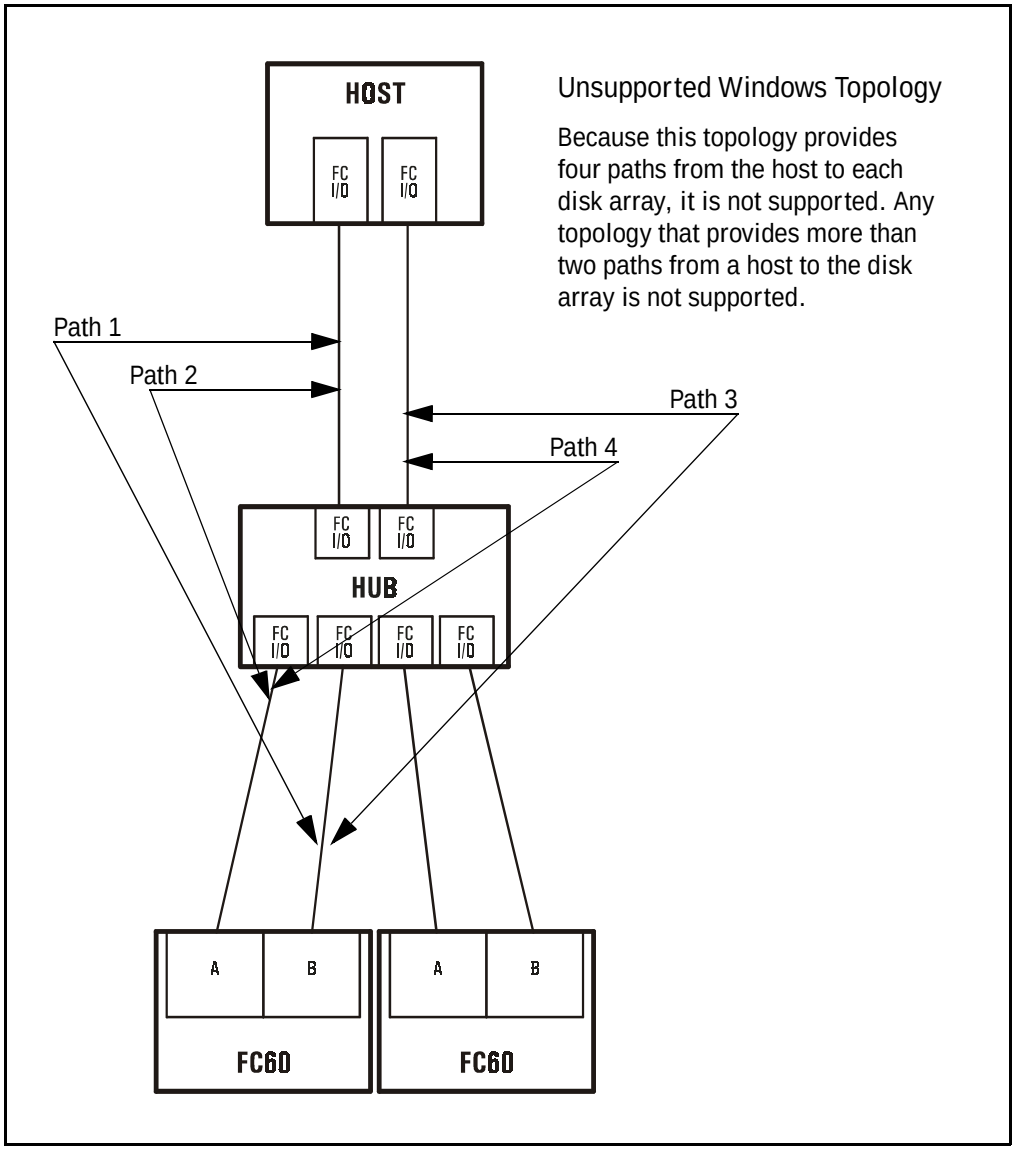



**Figure 44** An Unsupported Single-Hub Topology

## Non-High Availability Topologies

Figure 45 through Figure 47 illustrate non-high availability topologies. These topologies do not achieve the highest level of data availability because they have a hardware component that represents a single point of failure. That is, if the critical component fails, access to the data on the disk array will be interrupted. These topologies are simpler and less expensive to implement than true high availability topologies.

### Characteristics

- The critical components that represent a single point of failure in these configurations are:
  - The hub or switch.
  - The FC host adapter, if there is only a single connection from the host to the hub. In a cluster configuration a single host adapter does not represent a single point of failure because data will still be available through the other cluster host.
  - The FC cable if there is only a single connection between the host and the hub. In a cluster configuration the FC cable does not represent a single point of failure because data will still be available through the other cluster host.
- The use of hubs and switches increases the distance that can be achieved between the host and the Disk Array FC60. The host and hub can be up to 500 meters apart, and the distance from the hub to the disk arrays can be an additional 500 meters.

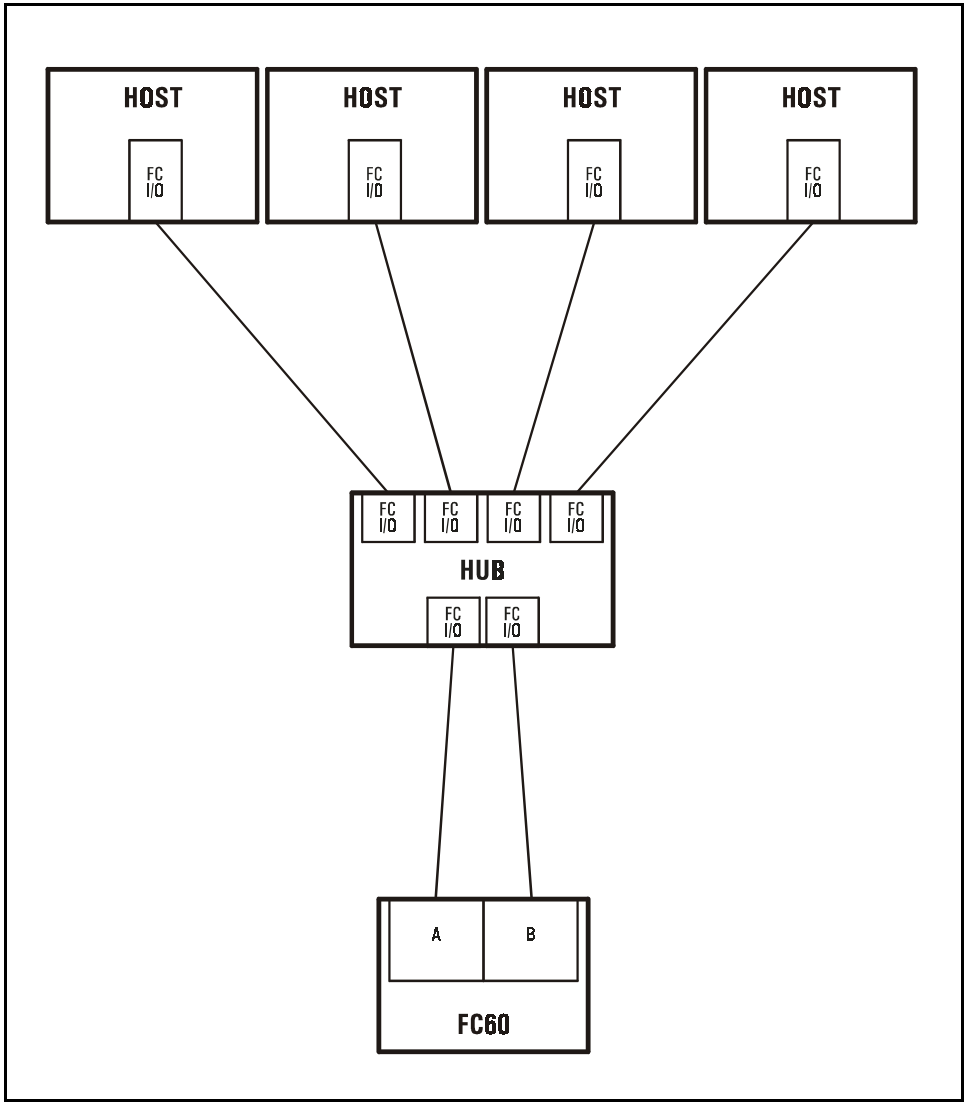



**Figure 45** Four Host/Single Hub/ Single Disk Array Non-HA Topology

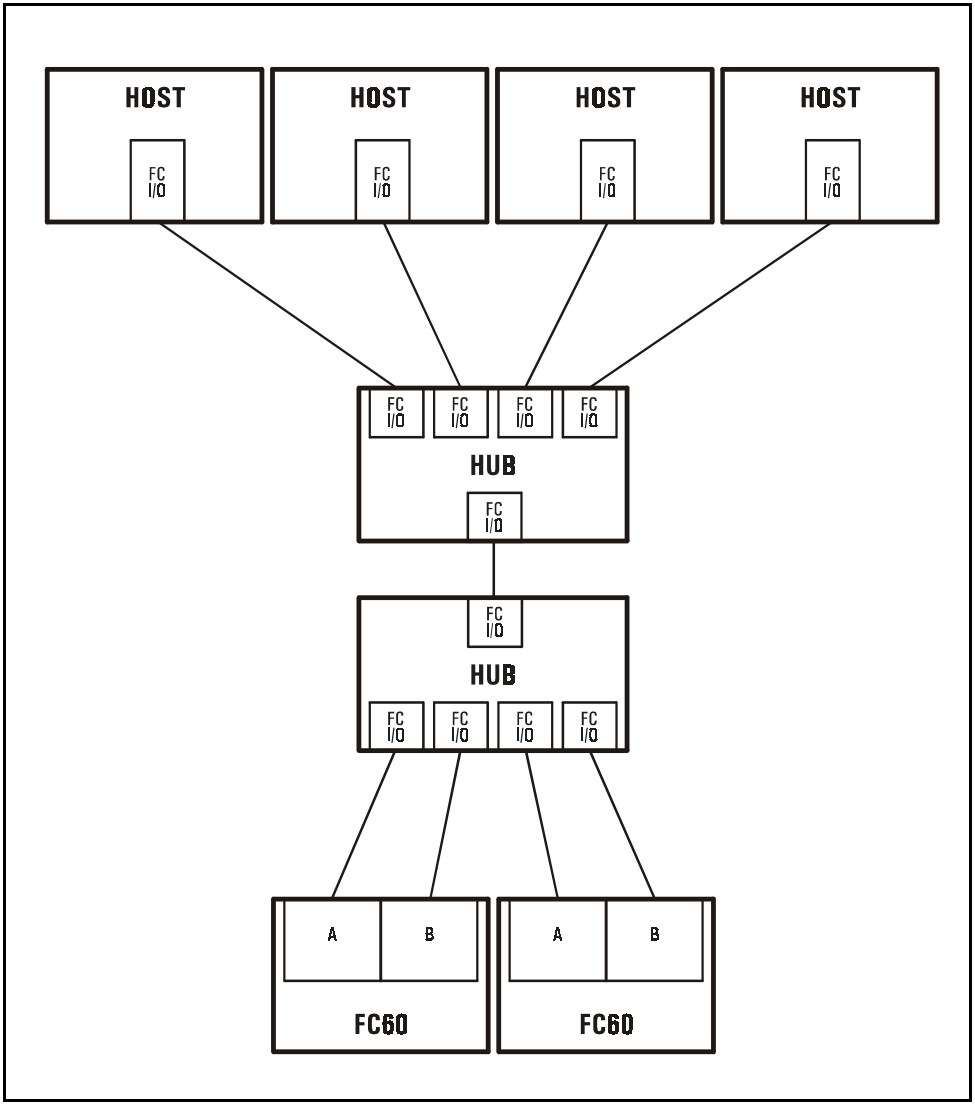



**Figure 46** Four Host/Cascaded Hubs/ Dual Disk Array Non-HA Topology

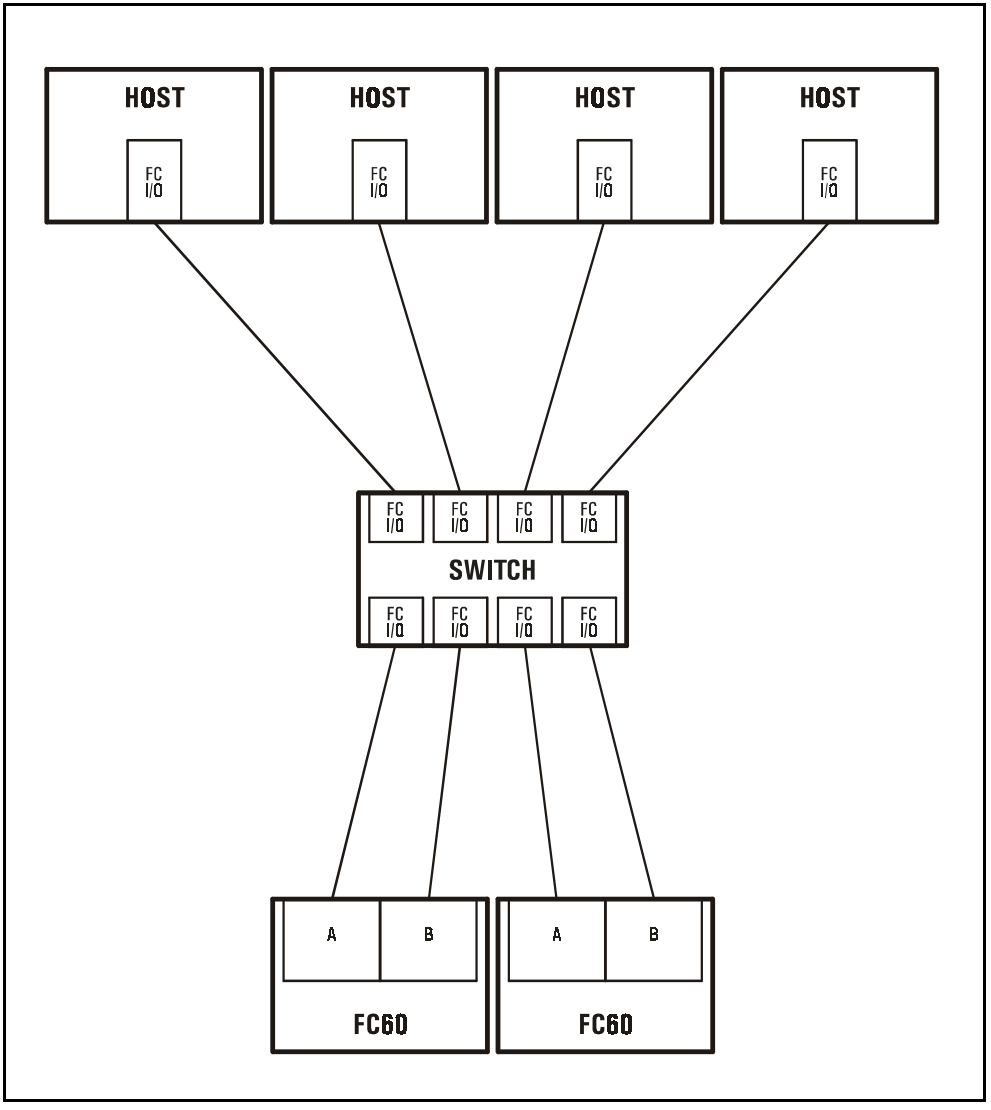



**Figure 47** Four Host/Single Switch/ Dual Disk Array Non-HA Topology

## High Availability Topologies

Figure 48 through Figure 51 illustrate high availability topologies. These topologies achieve the highest level of availability because they have fully redundant hardware data paths to each disk array. There are no single points of failure in these topologies. These topologies are more complex and expensive to implement than non-high availability topologies.

### Characteristics

- The direct connect topology (Figure 48) achieves high availability by providing redundant data paths between the host and the disk array, but this configuration also provides limited distance and connectivity.
- The use of multiple hubs creates redundant data paths between all hosts and all disk arrays. Multiple hubs also provides greater connectivity, increasing the number of disk arrays and host that can be interconnected.
- The use of hubs and switches increases the distance that can be achieved between the host and the Disk Array FC60. The host and hub can be up to 500 meters apart, and the distance from the hub to the disk arrays can be an additional 500 meters.

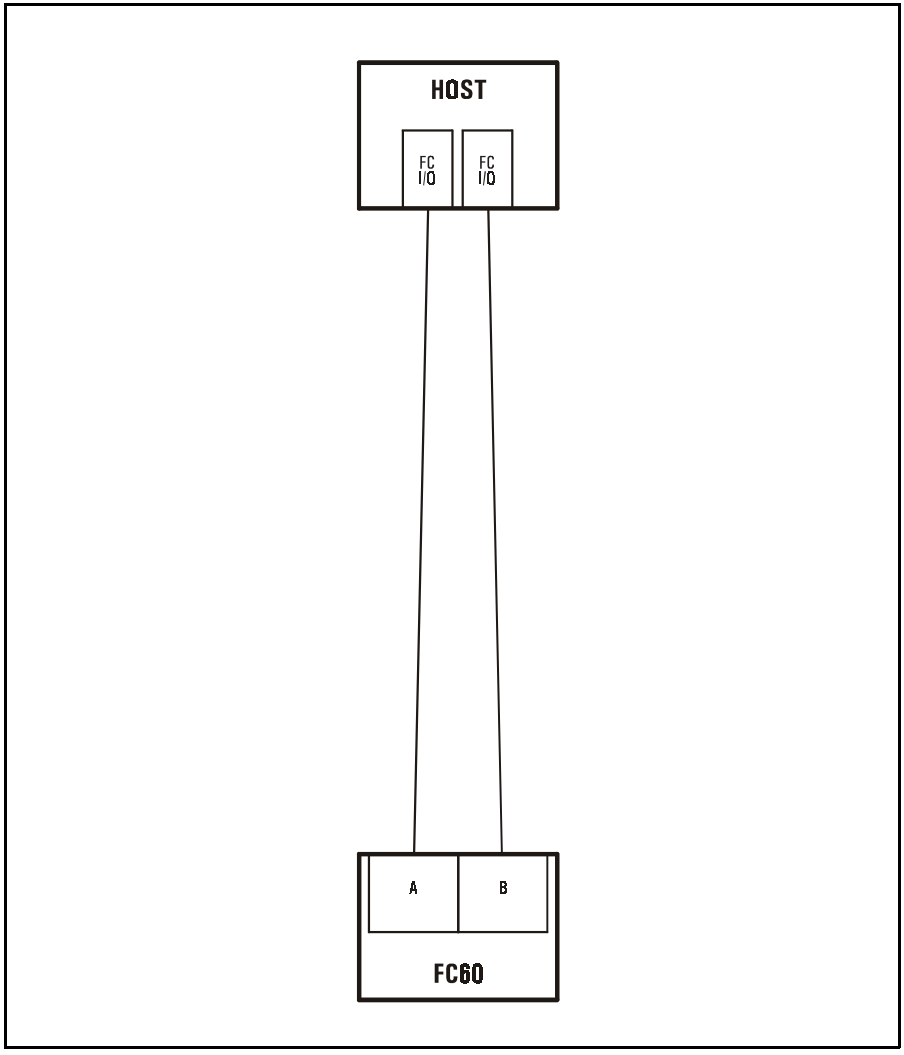



**Figure 48** Direct Connect Single Host/Single Disk Array HA Topology

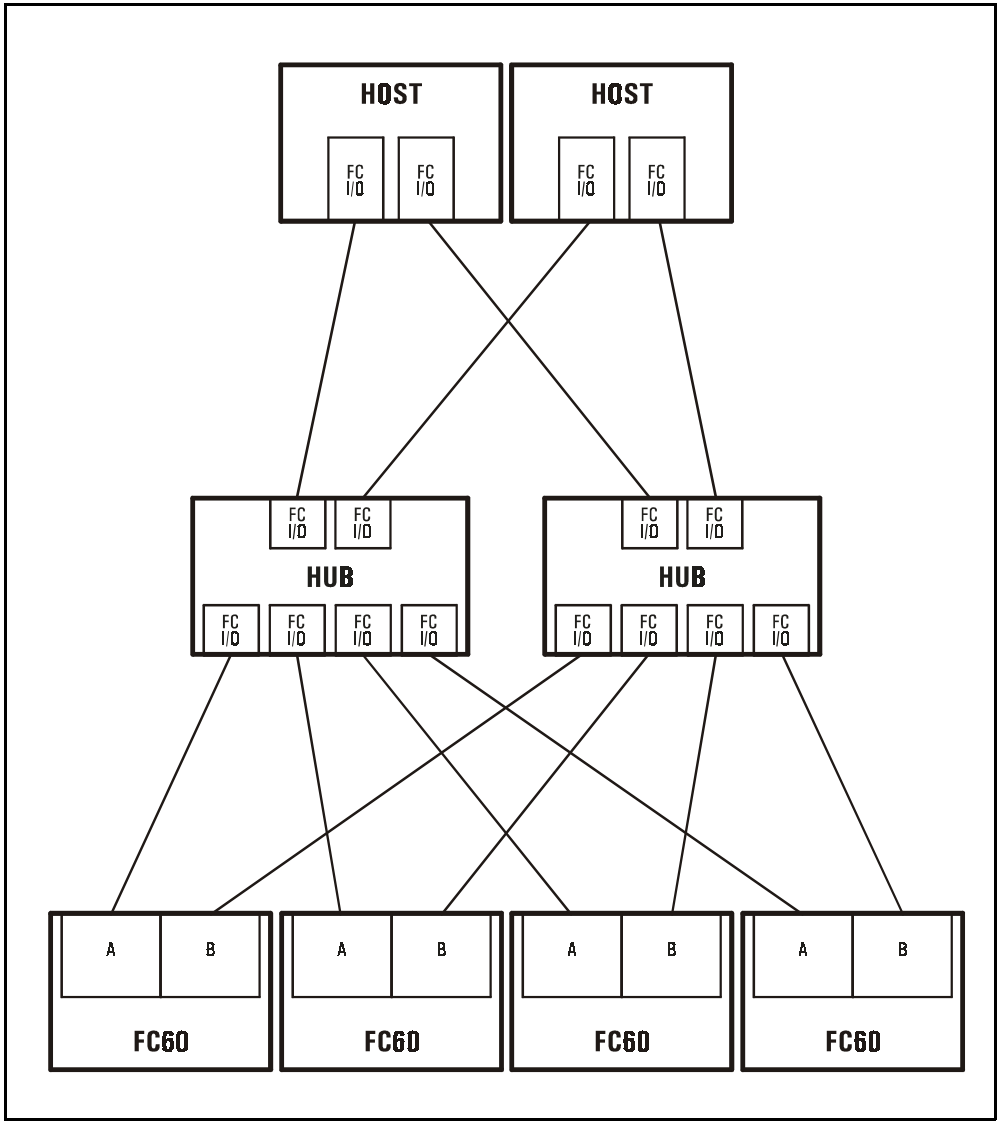



**Figure 49** Dual Host/Dual Hub/Four Disk Array HA Topology

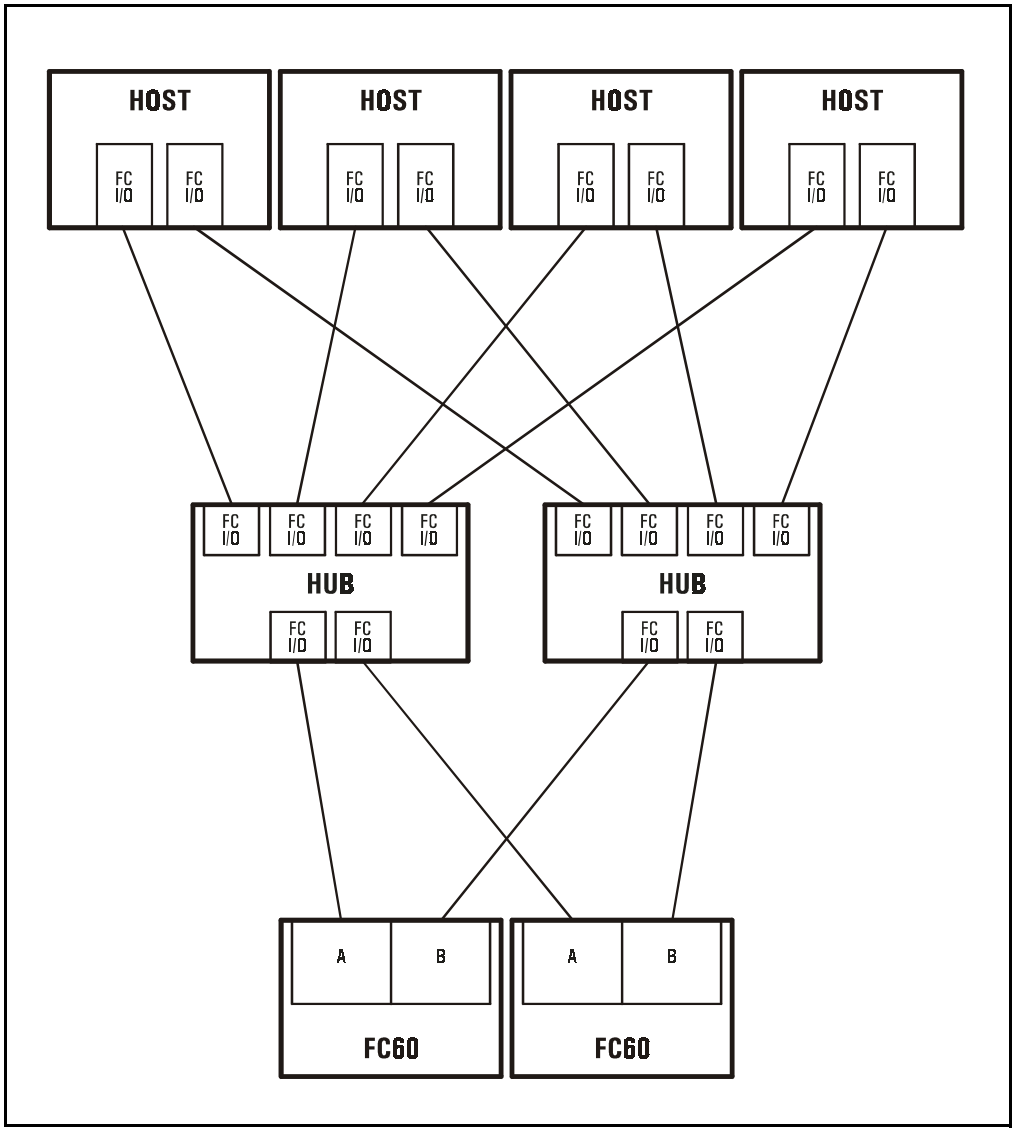



**Figure 50** Four Host/Dual Hub/Dual Disk Array HA Topology

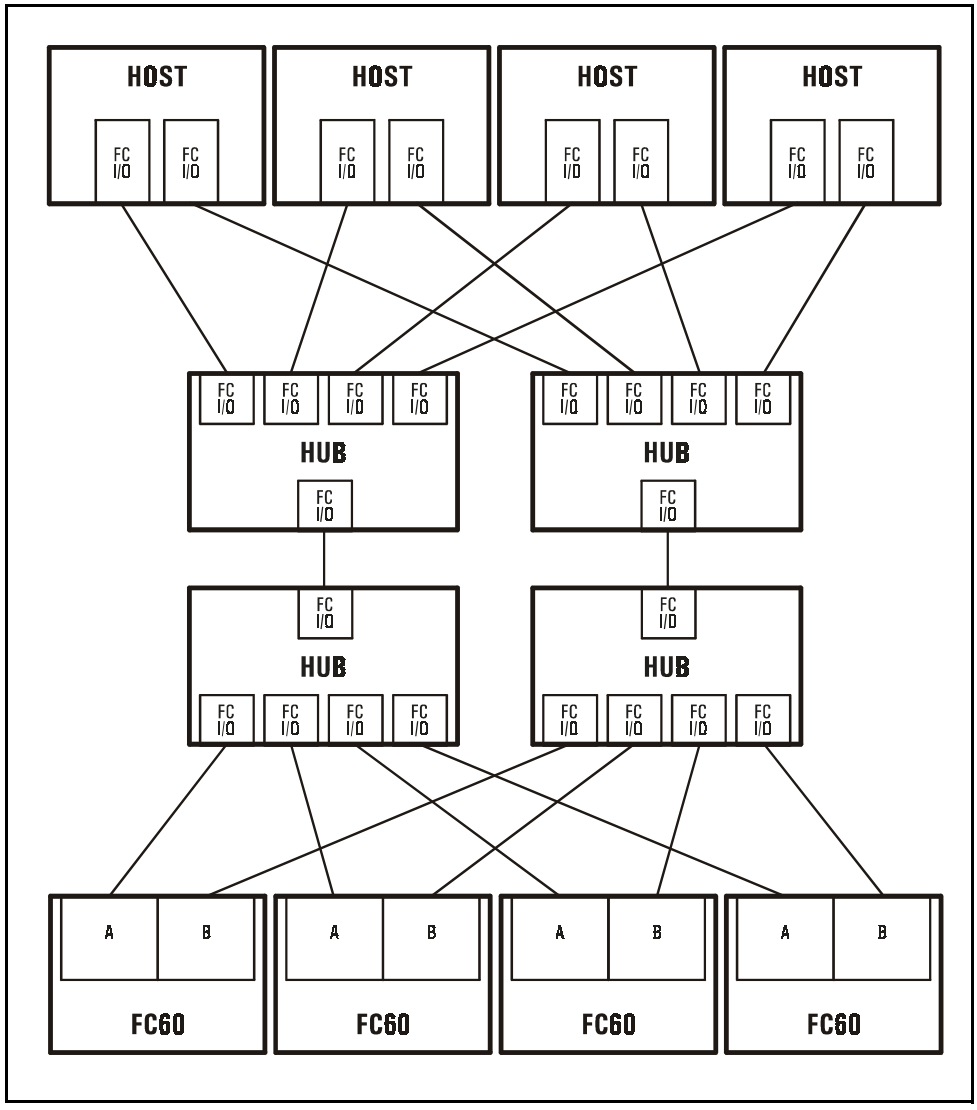



**Figure 51** Four Host/Dual Cascaded-Hubs/Four Disk Array HA Topology

# 3 INSTALLATION


Host System Requirements . . . . . 145
Site Requirements . . . . . 147
Power Distribution Units (PDU/PDRU) . . . . . 150
Installing the Disk Array FC60 . . . . . 155
Installing the Disk Enclosures . . . . . 160
Installing the Controller . . . . . 170
Configuration Switches . . . . . 176
Attaching Power Cords . . . . . 183
Attaching SCSI Cables and Configuring the Disk Enclosure Switches . . . . . 187
Connecting the Fibre Channel Cables . . . . . 196
Applying Power to the Disk Array . . . . . 198
Verifying Disk Array Connection . . . . . 206
Installing the Disk Array FC60 Software (HP-UX Only) . . . . . 213
Configuring the Disk Array . . . . . 216
Using the Disk Array FC60 as a Boot Device (HP-UX Only) . . . . . 222
Solving Common Installation Problems . . . . . 223
Adding Disk Enclosures to Increase Capacity . . . . . 224

# Overview

This chapter explains how to install the Disk Array FC60 enclosures into a cabinet and how to configure and connect the controller enclosure to the disk enclosures. It also covers the Fibre Cable connection to the host. Finally this chapter provides power up instructions and initial software installation requirements for operation of the disk array. Before installing the Disk Array FC60, the topology and array configuration should be established. See Chapter 2, Topology and Array Planning for information on these topics.

The Disk Array FC60 is designed to be installed in the following HP cabinets:

- Legacy (also referred to as original) HP cabinets (1.1m, 1.6m, and 2m)
- System/E Rack (which includes A490xA/J150xA: 1.1m, 1.6m, and 2m)

Installation consists of the following procedures:

- Preparing the site
- Installing PDUs
- Installing disk enclosures
- Installing the controller enclosure
- Setting configuration switches
- Attaching power cords
- Attaching SCSI cables and configuring the disk enclosure switches
- Connecting Fibre Channel cables
- Powering Up the Array
- Verifying Disk Array Connection
- Installing the Disk Array Management Software

# Host System Requirements

## HP-UX

The Disk Array FC60 is supported on the following host configurations:

- **Supported host platforms**: Table 13 lists the supported host platforms. The table also identifies which platforms support the disk array as a boot device when running HP-UX 11.x
- **Supported HP-UX versions**: 11.0, 11.11, and 10.20

**Note** The Disk Array FC60 is supported on HP-UX 11.11 with array firmware release HP03 only. Later versions of disk array firmware are not currently supported when running HP-UX 11.11.

- **Software requirements**: The Disk Array FC60 management software must be installed on the host. This software is distributed on *HP-UX Support Plus CD-ROM*, release 9909 or later.



**Table 13 Supported Host Platform Information**

| Supported Host | Boot Support on HP-UX 11.x? | Fibre Channel I/O Adapter |
|---|---|---|
| K-class | Yes | A3404A |
| V-class | Yes | A5158A, A3740A |
| L-class | Yes | A5158A, A3740A |
| D-class | Yes | A3591B |
| N-class | Yes | A5158A, A3740A |
| R-class | Yes | A3591B |
| T-class | No | A3636A |

**Table 13 Supported Host Platform Information (cont'd)**

| Supported Host | Boot Support on HP-UX 11.x? | Fibre Channel I/O Adapter |
|---|---|---|
| C-class | No | A5158A on HP-UX 11.x<br>A3740A on HP-UX 10.20 |
| A4xx-A5xx class | Yes | A5158A |

### Fibre Channel I/O Adapters

The host must have the correct adapter installed. The supported host adapters are listed in Table 13. Fibre Channel drivers are provided with the operating system. For complete information on Fibre Channel I/O adapter installation and operation and driver installation, refer to the *Fibre Channel Mass Storage Adapters Service and User Manual (*part number A3636-90002). A copy of the manual may be viewed or printed by accessing the following web site:

www.hp.com/essd/efc/A3636A_documentation.html

## Windows NT and Windows 2000

For information on Windows NT and Windows 2000 host system requirements, see the *HP Storage Manager 60 User's Guide* included with the *HP Storage Manager 60* software (A5628A).

# Site Requirements

## Environmental Requirements

The area around the array must be cooled sufficiently so it does not overheat. Chapter 8, Reference and Regulatory, contains environmental specifications for the Disk Array FC60. Refer to that section for the required environmental specifications.

## Electrical Requirements

The site must be able to provide sufficient power to meet the needs of the devices in the cabinet(s). For high availability operation, two separate sources, preferably an uninterruptible power source (UPS), must to be provided. Make sure that the UPS or other power source is capable of supplying an amount of power equal to the sum of the VA units required to power the Disk Array FC60 and any other peripherals installed in the cabinet. For more information about racking configurations see “Cabinets and Uninterruptible Power Supplies” in the current *HP 9000 Servers Configuration Guide*.

The Disk Array FC60 is scalable and may contain from one to six disk enclosures. The power requirement varies according to the number of disk enclosures installed. Providing power for a full six-disk enclosure will ensure sufficient power in the event of future scaleup.

Table 14 shows the steady state operating current and in-rush current for the controller options. Use this table to determine the proper number of PDUs or PDRUs and number of receptacles you will need. Make sure all electrical wiring to the service point (plug) is adequately sized to carry in-rush and steady state currents.

**Table 14** Total Operating and In-Rush Currents

| | Operating Current @ 110v | Operating Current @ 220v | In-Rush* Current | Power Cords |
|---|---|---|---|---|
| Controller w/ 6 Disk Enclosures | 41.3A | 20.4A | 124A | 14 |
| Controller w/ 5 Disk Enclosures | 34.8A | 17.2A | 104A | 12 |
| Controller w/ 4 Disk Enclosures | 28.3A | 14.0A | 84A | 10 |
| Controller w/ 3 Disk Enclosures | 21.8A | 10.8A | 64A | 8 |
| Controller w/ 2 Disk Enclosures | 15.3A | 7.6 | 44A | 6 |
| Controller w/ 1 Disk Enclosure | 8.8 | 4.4A | 24A | 4 |

* In-rush current occurs for 10 to 12 milliseconds

**Table 15** Controller Enclosure Electrical Requirements

| Measurement | Value @ 110v | Value @ 220v |
|---|---|---|
| Voltage<br>– Range<br>– Frequency | <br>100-127V<br>50 - 60 Hz | <br>220-240V<br>50 - 60 Hz |
| Current<br>– Typical<br>– Maximum Operating<br>– Inrush | <br>1.5A<br>2.3A<br>21.4A | <br>0.8A<br>1.2A<br>42.9A |

**Table 16** Disk Enclosure Electrical Requirements

| Measurement | Value |
|---|---|
| Voltage<br>– Range<br>– Frequency | <br>220-240V 100 - 127V<br>50-60Hz |
| Current<br>– Typical<br>Maximum Operating<br>– 100 - 120 V<br>– 200 - 240 V<br>– Maximum In-rush | <br>2.9 - 3.2A<br><br>2.6 - 3.2A<br>5.3 - 6.7A<br>20A |

* In-rush current occurs for 10 to 12 milliseconds

HP recommends the use of magnetic-type circuit breakers, which are capable of handling large in-rush currents for short durations (10 to 12 milliseconds) and are rated adequately for the steady state currents.

**Table 17** Recommended European Circuit Breakers*

| Number of Components | Breaker Rating | Breaker Type |
|---|---|---|
| 1 to 3 | 16 amps | Type C or Type D per IEC 898<br>**or**<br>Type K per IEC 947-2 |
| 4 | 16 amps | Type D per IEC 898<br>**or**<br>Type K per IEC 947-2 |

* Data assumes no other devices share the circuit breaker.

**Note** Circuit breaker rating must be adequate for the total current drawn by all devices on all electrical paths that share a circuit breaker.

# Power Distribution Units (PDU/PDRU)

PDUs provide a sufficient number of receptacles for a large number of devices installed in a cabinet. A PDU connects to the site power source and distributes power to its ten receptacles. The disk array power cords are connected to the PDUs and each PDU is connected to a separate UPS. For high availability, two PDUs should be installed and connected to separate site power sources.

PDUs come in a variety of heights, from five feet down to 19 inches. Because the 19-inch PDU is shorter, it can be installed both vertically and horizontally. The peak power requirements and PDU capacity affect the number of electrical devices that can be installed in a rack. For example, to install more than four disk enclosures in an HP legacy rack, you must upgrade from the older 3-foot and 5-foot PDUs to 19-inch PDUs.

In addition to rack density, consider the following in choosing PDUs:

- **Redundant power source.** To connect redundant power supplies to separate PDUs, install redundant PDUs.
- **Number of cords to the AC source**. Using 30-amp PDRUs instead of 16-amp PDUs reduces the number of cords to the wall.
- **Future needs**. Installing surplus PDU capacity allows you to add devices later.
- **In-rush margins.** For installations that require four or more 16-amp PDUs, HP recommends HP 30-amp PDRUs (E7681A, E7682A) for their inherent in-rush protection.
- **On/Off switch support**. Some PDU/PDRU options support the use of a single-point on/off switch. See Table 18 and Table 19.

The following tables show recommended PDU/PDRU combinations for one or more components in a rack. Data assumes 220V AC nominal power and redundant PDU/PDRUs. For nonredundant configurations, divide the number of recommended PDU/PDRUs by 2.

**Table 18** Recommended PDU/PDRUs for HP Legacy Racks

| Number of Components | 1.1-meter (21 U) Rack | 1.6-meter (32 U) Rack | 2.0-meter (41 U) Rack |
|---|---|---|---|
| 1 – 4 | Two 3-foot/16-amp PDUs*<br>***or***<br>Two 19-inch/16-amp PDUs | Two 5-foot/16-amp PDUs*<br>***or***<br>Two 19-inch/16-amp PDUs | |
| 5 – 8 | NA** | Two 19-inch/30-amp PDRUs | |
| 9 – 10 | NA** | NA** | Four 19-inch/30-amp PDRUs |

* Supports cabinet on/off switch.
** Rack height does not allow additional enclosures.

**Table 19** Recommended PDU/PDRUs for HP System/E Racks

| Number of Enclosures | 1.25-meter (25 U) Rack | 1.6-meter (33 U) Rack | 2.0-meter (41 U) Rack |
|---|---|---|---|
| 1 – 4 | Two 19-inch/16-amp PDUs ***or*** Two 19-inch/30-amp PDRUs* | | |
| 5 – 8 | NA** | Two 19-inch/30-amp PDRUs* | |
| 9 – 11 | NA** | NA** | 4 19-inch/30-amp PDRUs |

* Supports the cabinet on/off switch option.
** Rack height does not allow additional enclosures.

## Installing PDUs

Choose PDU/PDRU locations with the following guidelines in mind:

- Place PDU/PDRUs within the reach of power cords.
- Place PDUs vertically whenever possible. See sample installations in Figure 52 and Figure 53. Installing PDUs horizontally can interfere with accessibility to units behind the PDU. (PDRUs must be installed vertically, as per the installation instructions.)
- Place vertical PDU/PDRUs on each side of the cabinet so that the two power cords from the enclosure's redundant power supplies do not have to cross over replaceable components in the middle of the product.
- Thirty-amp PDRUs must be installed directly behind the disk enclosure to achieve the highest densities in 2-meter racks. Hinged brackets allow HP's 30-amp PDRUs (HP E7681A and E7682A) to swing aside for servicing the disk enclosures (see Figure 53).

Refer to the documentation supplied with the PDU/PDRU for installation instructions.

## Recommended UPS Models

The following Hewlett-Packard Power Trust models are recommended for use with the HP SureStore E Disk Array FC60. Each UPS supplies up to 15 minutes of standby power.

**Table 20** Recommended UPSs

| Disk Array Configuration | UPS Product Number | Output Power | Package |
|---|---|---|---|
| Controller enclosure and up to two disk enclosures | A2998B | 3.0 kVA (3.0 kW) | Rackmount |
| Controller enclosure and up to four disk enclosures | A3589B | 5.5 kVA (5.5 kW) | Rackmount |
| Controller enclosure and up to six disk enclosures | A1350A | 8.0 kVA (8.0 kW) | Standalone |

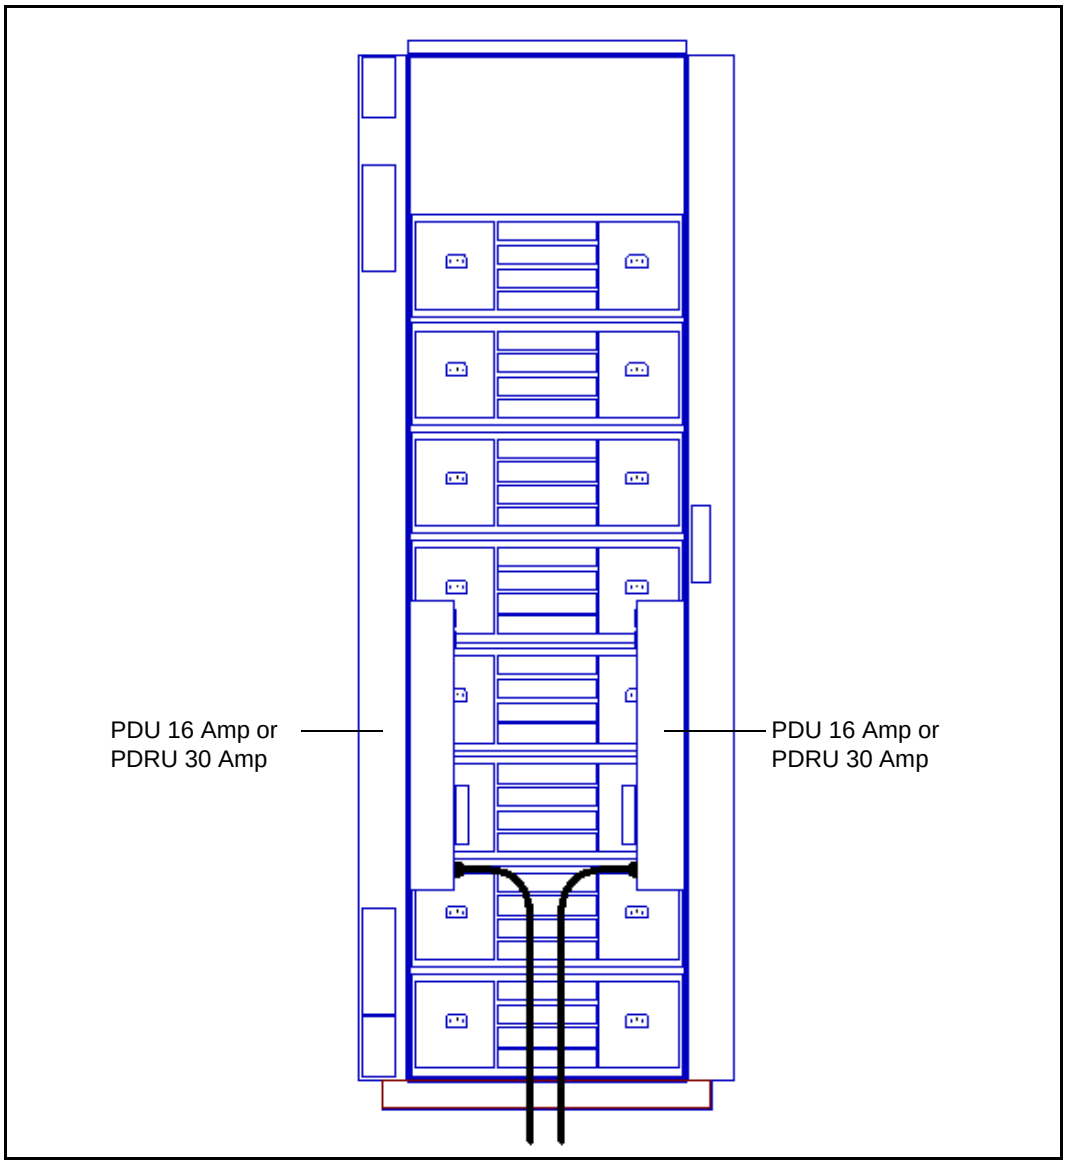



**Figure 52** PDU Placement in 1.6-Meter Rack

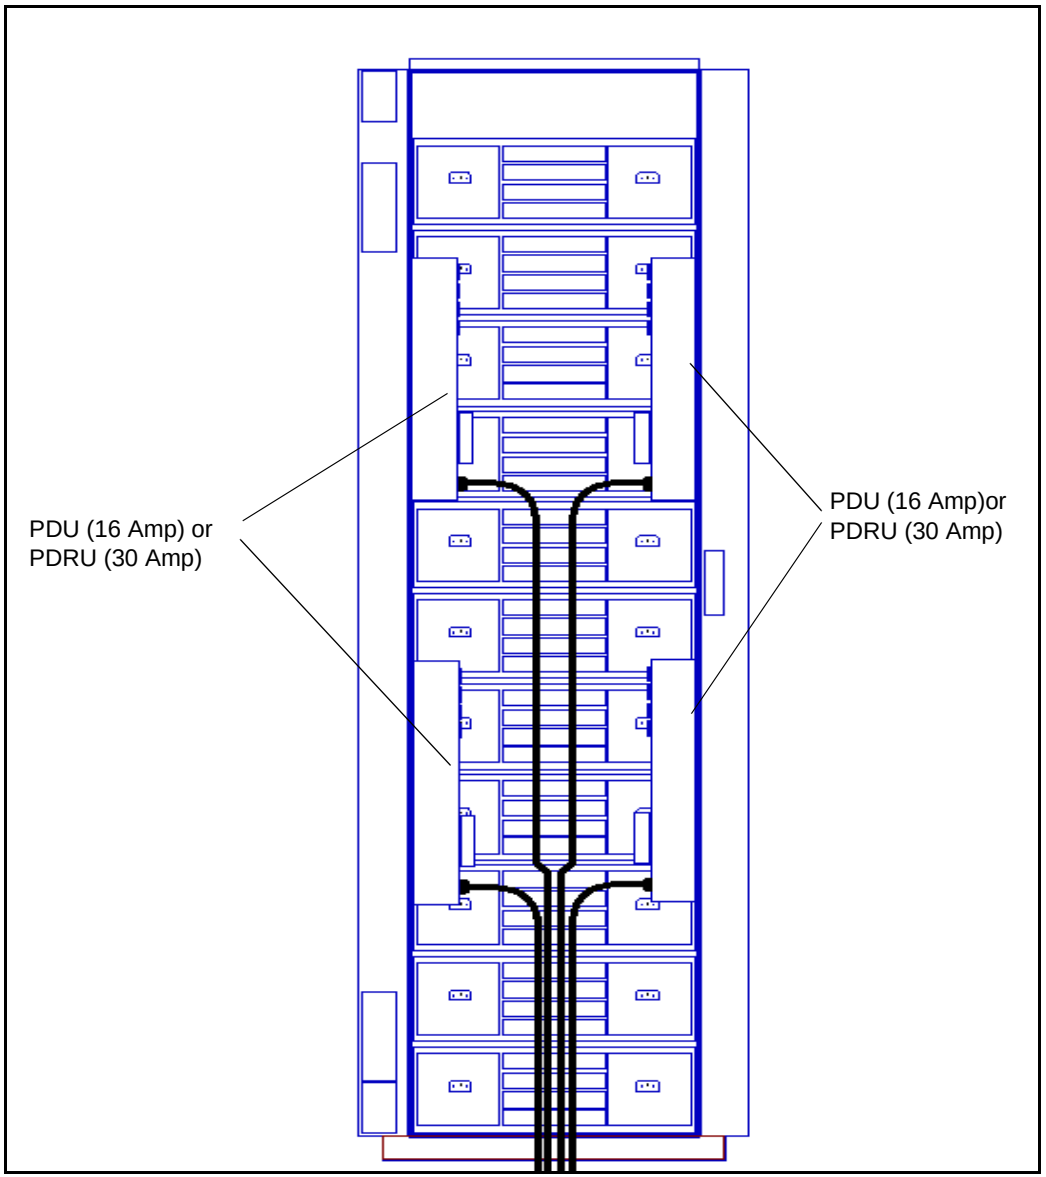



**Figure 53** PDRU Placement in 2.0-Meter Rack

# Installing the Disk Array FC60

**Note** The A5277AZ factory assembled disk array is fully configured and requires only connection to the host. Proceed to "Connecting the Fibre Channel Cables" on page 196 to complete the installation. Do not remove any of the factory installed disk enclosures before powering on the disk array, or the disk array will not function properly.

If the factory default Fibre Channel address settings (4 and 5) will not work, the settings can be changed as described in "Fibre Channel Host ID Address Setting" on page 179

Installation of the Disk Array FC60 consists of installing the controller enclosure and up to six disk enclosures. Before performing this installation you should have determined the configuration for the array (as described in chapter 2, Array Planning and Topology). You should also review this section prior to performing the installation.

The Disk Array FC60 is supported in the System/E racks and in the HP legacy (original) cabinets.

**Note** The HP SureStore E Disk Array FC60 has been tested for proper operation in supported Hewlett-Packard cabinets. If the disk array is installed in an untested rack configuration, care must be taken to ensure that all necessary environmental requirements are met. This includes power, airflow, temperature, and humidity. Failure to meet the required operating specifications may result in product failure.

**Table 21** EIA Spacing for Racks and Array Enclosures

| Component | Measure (EIA Units) |
|---|---|
| **Legacy Cabinets (1 EIA Unit = 1.75")** | |
| 1.1 Meter Cabinet | 21 EIA Units, Total Available |
| 1.6 Meter Cabinet | 32 EIA Units, Total Available |
| 2.0 Meter Cabinet | 41 EIA Units, Total Available |
| Controller Enclosure FC60 | 5 EIA Units Used (includes 1/2 rail space below and remaining 1/2 EIA unit above enclosure) |
| Disk Enclosure SC10 | 4 EIA Units Used (3.5 disk enclosure plus 1/2 rail space) |
| **System/E Racks (1 EIA Unit = 1.75")** | |
| 1.1 Meter | 25 EIA Units, Total Available |
| 1.6 Meter | 33 EIA Units, Total Available |
| 2.0 Meter | 41 EIA Units, Total Available |
| Controller Enclosures FC60 | 4 EIA Units Used |
| Disk Enclosures SC10 | 3.5 EIA Units Used |

**Note** IMPORTANT: When installing an enclosure in a rack, if there is a space between enclosures that space must be filled with a filler panel to maintain proper air flow. This is a requirement for proper cooling.

When installing the array in the cabinet, install the enclosures from the bottom up (to prevent the cabinet from becoming top heavy and avoid a potential tipping hazard). Refer to the document supplied with the rack for proper installation information. Typically, disk enclosures are installed from the bottom of the rack up, with the controller enclosure installed directly above the disk enclosures.

Figure 54 shows EIA rack locations for installation of six disk enclosures and one controller enclosure FC60 in a System/E rack. Special rails for the System/E racks allow installation of enclosures at either full EIA or mid-EIA (1/2-EIA) position. This allows the

installation to utilize 1/2 EIA units available from the disk system SC10's 3.5 EIA unit height.

Figure 55 shows rack locations for installation of six disk enclosures and one controller enclosure (positioned on top) for legacy racks. When disk enclosures are installed in legacy racks, an unusable 1/2-EIA space is left at the bottom of the enclosures. This space must be filled with a 1/2-EIA unit filler for each enclosure installed.

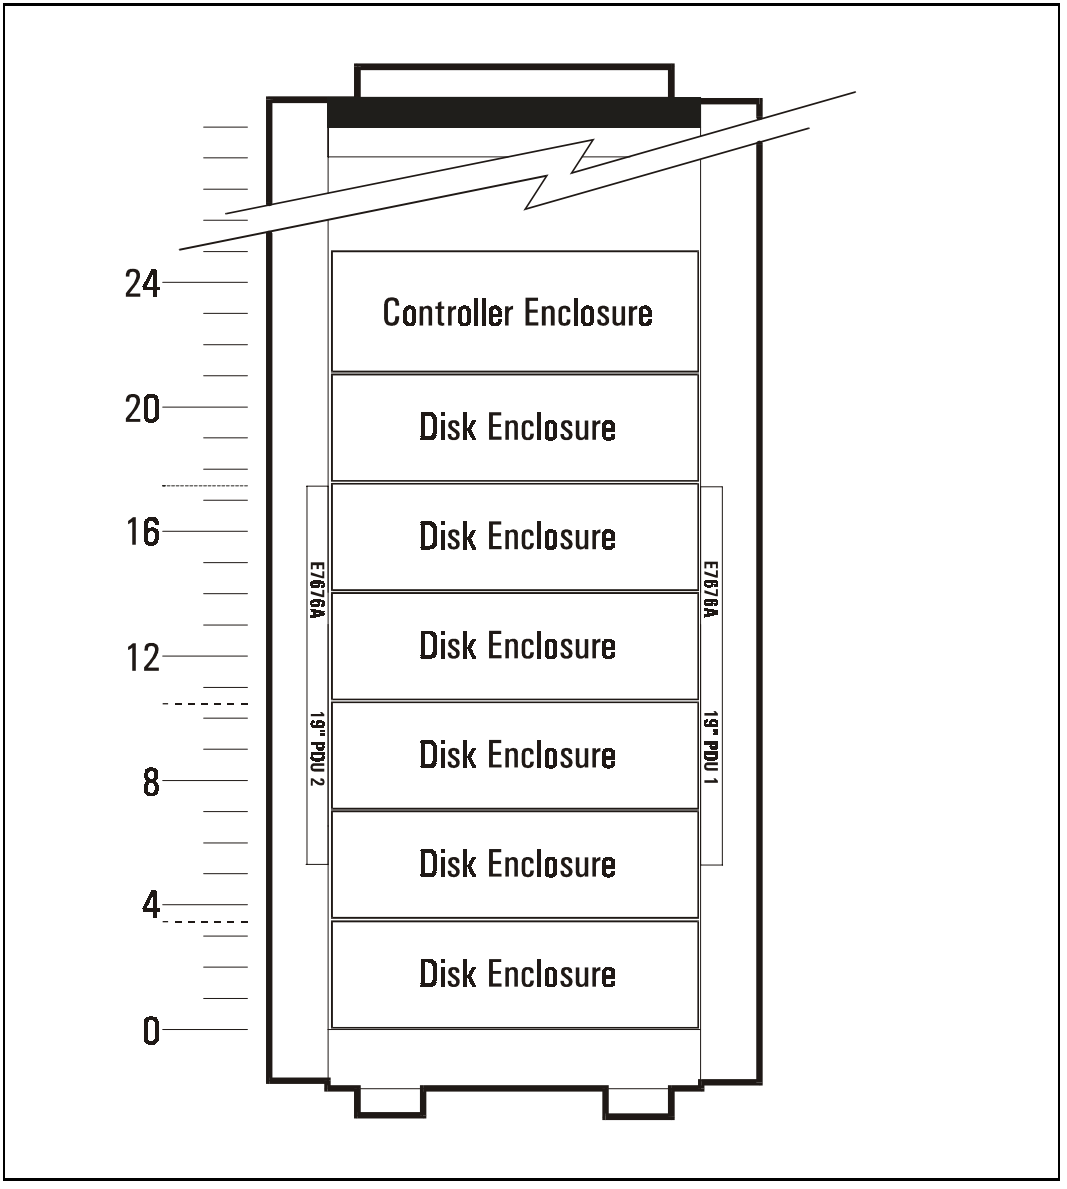



**Figure 54** Enclosure EIA Positions for System/E Racks

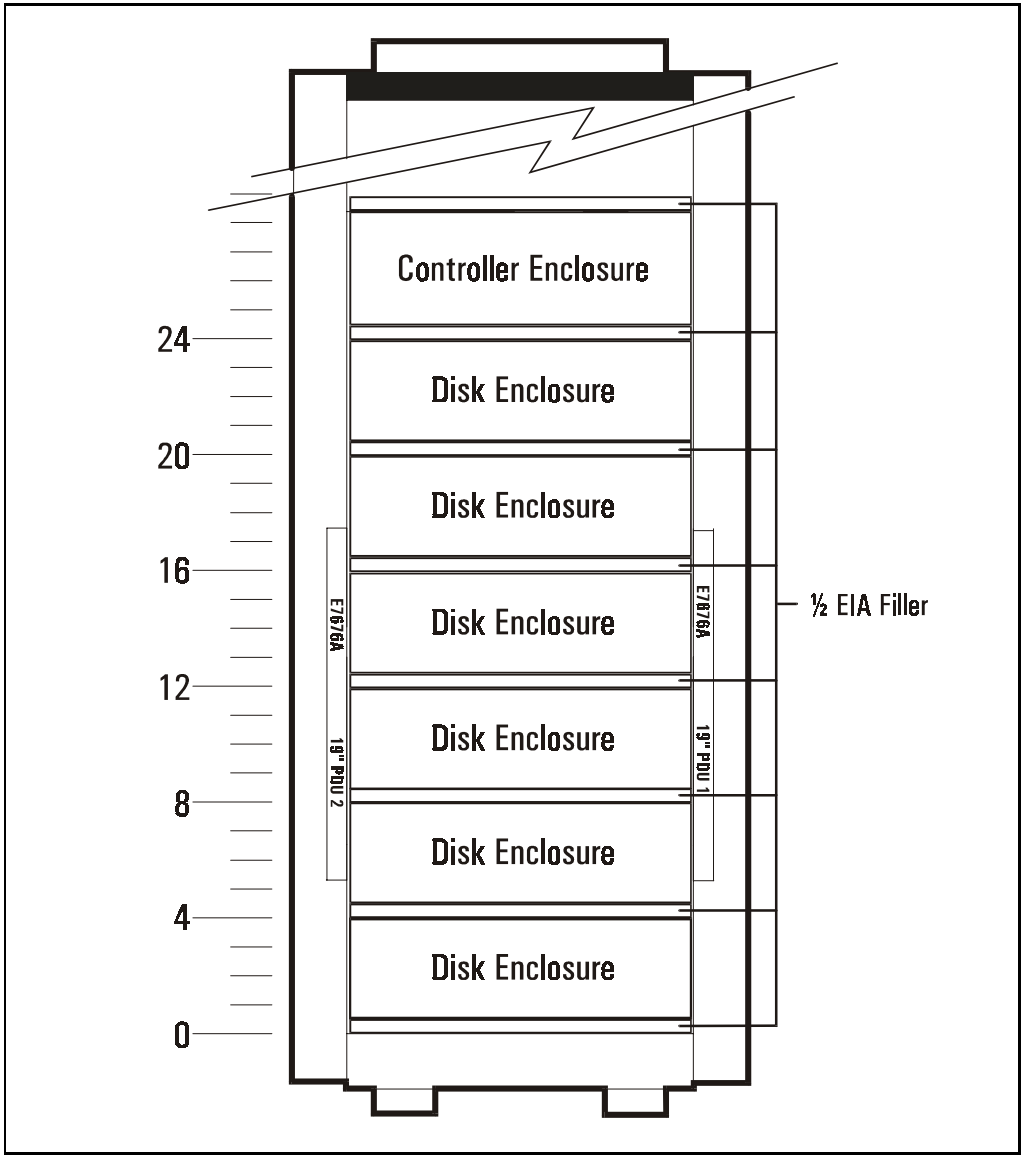



**Figure 55** Enclosure EIA Positions for Legacy Cabinets

# Installing the Disk Enclosures

Disk enclosures should be installed in the rack starting at the bottom and proceeding upward. When all disk enclosures are installed, the controller enclosure is installed at the top, directly above the disk enclosure. Installation instructions for the disk enclosure SC10 are provided below; installation instructions for the controller enclosure FC60 follow this section.

## Step 1: Collect Required Tools

- Torx T25 screwdriver
- Torx T15 screwdriver
- Small flat-blade screwdriver

## Step 2: Unpack the Product

1. Lift off the overcarton and verify the contents of the box (see Table 22 and Figure 56).

**Table 22** Disk Enclosure SC10 Accessories

| Figure Label | Part (part number) |
|---|---|
| A | Disk Enclosure (with power supplies, fans, and BCC modules) |
| B | Rail kit (A5251A) for System/E cabinets |
| C | Rail kit (A5250A) for legacy cabinets, 1 ea. |
| D | Rack filler panel, 1/2 EIA Unit, 1 ea. |
| E | Disks or fillers (A5234A or A5235A), pack of 10 |
| F | SCSI terminators, 2 ea. |
| G | Power cords, 2ea |
| H | ESD strap (9300-2170), 1 ea. |
| I | Documentation |

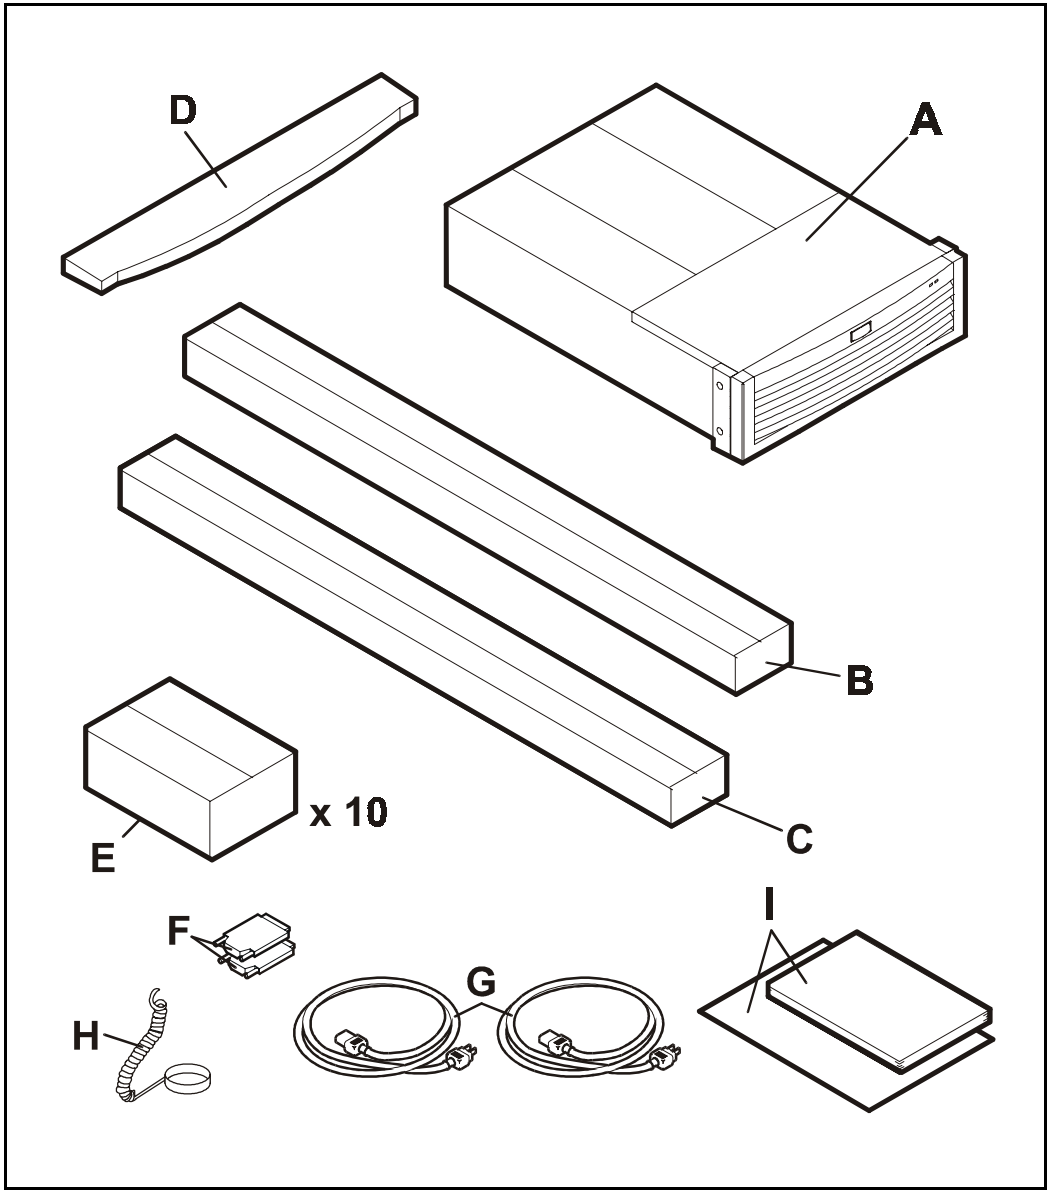



**Figure 56** Disk Enclosure Contents

## Step 3: Install Mounting Rails

Select the rail kit for the appropriate rack and follow the instructions included with the rail kit to install the rails in the rack. The following rail kits are available for use with the disk enclosure:

- HP A5250A for legacy HP Racks (C2785A, C2786A, C2787A A1896A, or A1897A)
- HP A5251A for HP Rack System/E
- HP 5656A for Rittal 9000 racks

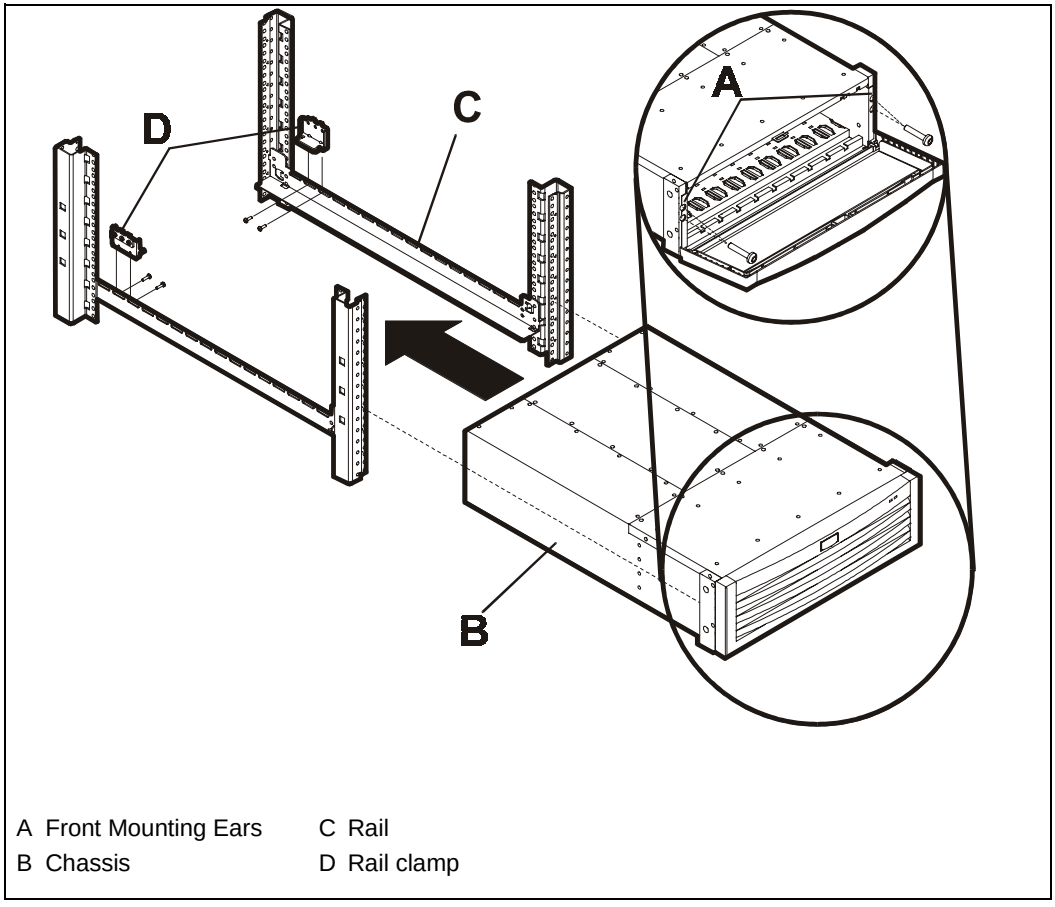
## Step 4: Install the Disk Enclosure

***CAUTION*** *Do not try to lift the disk enclosure using the handles on the power supplies.*

1. (Optional) Remove the power supplies to prepare the disk enclosure for lifting:
   a. With the chassis still in the box, loosen the screw in the handle of each power supply.
   b. Pull the cam handle down to disengage the power supply from the backplane and pull each power supply out of the chassis. Support the far end of the supply with your free hand as it clears the chassis.
   c. Set the power supplies aside to be reinstalled later.

**WARNING** **Do not attempt to lift the disk enclosure without the help of another person or a lift device. Even without power supplies and disk drives, it weighs 50 pounds.**

2. With another person or using a lift device, lift the disk enclosure and slide it, back-end first, into the front of the rack and onto the previously installed rails (Figure 57). Push the enclosure as far into the rack as it will go.




A Front Mounting Ears
B Chassis
C Rail
D Rail clamp

**Figure 57** Mounting the Disk Enclosure (Rack System/E shown)

***CAUTION*** *To protect the door, do not lift or move the disk enclosure with the door open.*

3. Unlock and open the disk enclosure door, using a thin flat-blade screwdriver to turn the lock (Figure 58).

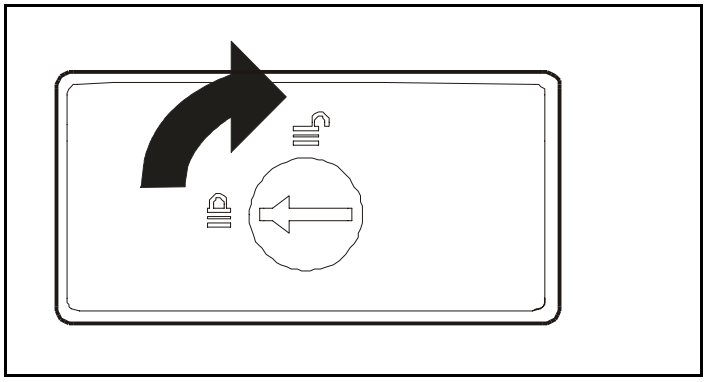


**Figure 58** Door Lock

4. Ensure that one hole in each mounting ear (A in Figure 57) aligns with the sheet metal nuts previously installed on the rack front columns.
5. Insert two screws (A in Figure 57) through the matching holes in the disk enclosure mounting ears and rack front columns. Tighten screws.
6. Close the door.

7. If using an HP rack, fasten the back of the disk enclosure to the rails using the rear hold-down clamps from the rail kit.

   a. If you are installing the disk enclosure in an HP legacy rack, set the clamp (A in Figure 59) on top of the rail (B) so that the tabs point up and the screw holes are on the slotted side of the rail. Skip to step c.

   b. If you are installing the disk enclosure in an HP Rack System/E, set the clamp (D in Figure 57)

   c. Push the clamp tight against the back of the disk enclosure. The raised tab of the clamp should overlap the bottom edge of the disk enclosure chassis.

   d. Insert and tighten two M5 16mm screws through each clamp and rail.

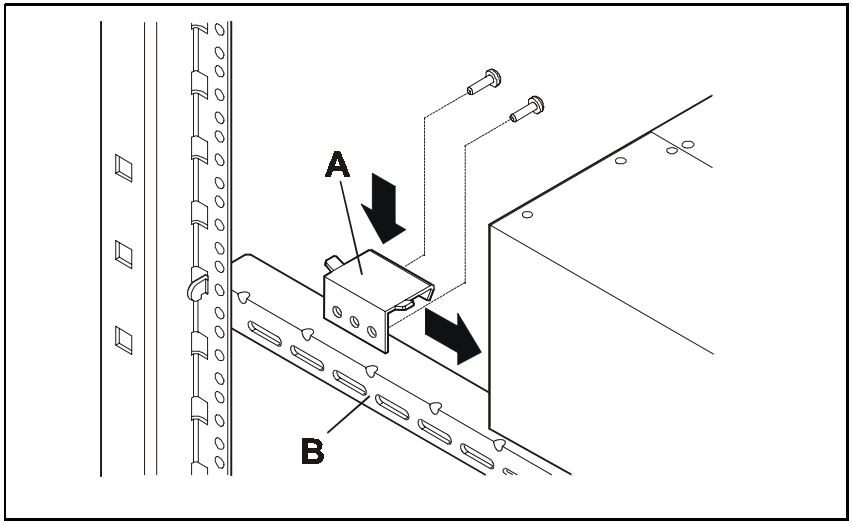



**Figure 59** Legacy Rack Rail Clamp Installation, Disk Enclosure

8. Reinstall power supplies removed in Step 1.

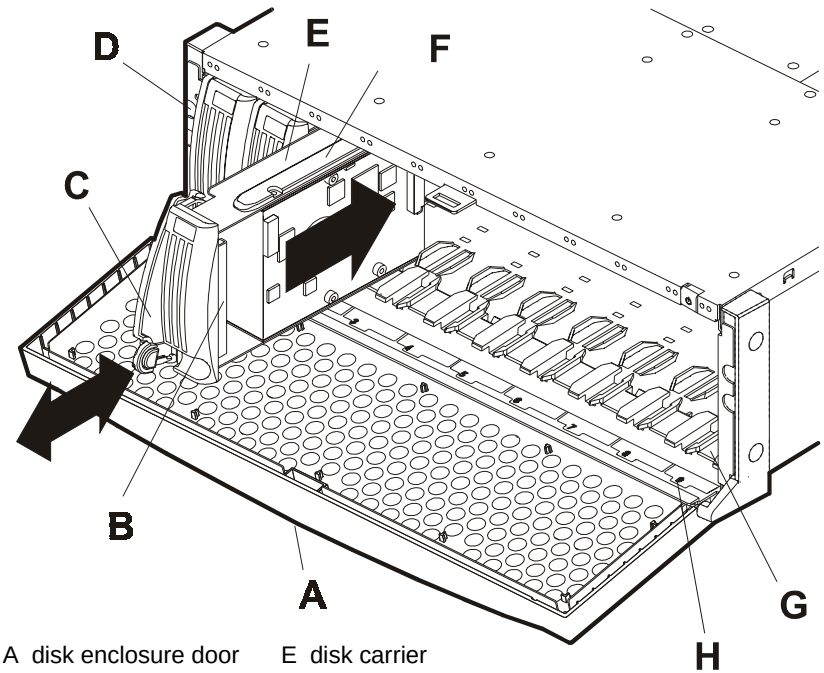
9. Install half-EIA unit rack filler panel(s) as needed. A half-EIA unit gap exists between products in legacy HP racks and when an odd number of disk enclosures are installed in the Rack System/E.

## Step 5: Install Disks and Fillers

***CAUTION*** ***Touching exposed areas on the disk can cause electrical discharge and disable the disk. Be sure you are grounded and be careful not to touch exposed circuits.***



Disks are fragile and ESD sensitive. Dropping one end of the disk two inches is enough to cause permanent damage. Static electricity can destroy the magnetic properties of recording surfaces. Grip disks only by their handles (B in Figure 60) and carriers, and follow strict ESD procedures.




A disk enclosure door
B disk handle
C cam latch
D ESD plug
E disk carrier
F insertion guide
G slot guide
H slot number

**Figure 60** Disk Installation

1. Open the disk enclosure door.
2. Put on the ESD strap (provided with the accessories) and insert the end into the ESD plug-in (D in Figure 60) near the upper left corner of the disk enclosure.

***CAUTION*** *Disks are fragile. Handle them carefully.*

3. Remove the bagged disk from the disk pack.

***CAUTION*** ***Do not touch exposed circuit board side of the disk module.***

4. Remove the disk from the ESD bag, grasping the disk by its handle (B).
5. Using a pencil, mark an “X” on the Slot Location Map label (A in Figure 61), located on the top of the disk module. This mark identifies the disk slot the disk module should be installed in case it is removed. It is important to return a disk to its same location to maintain disk module addressing used by some FC60 software applications.

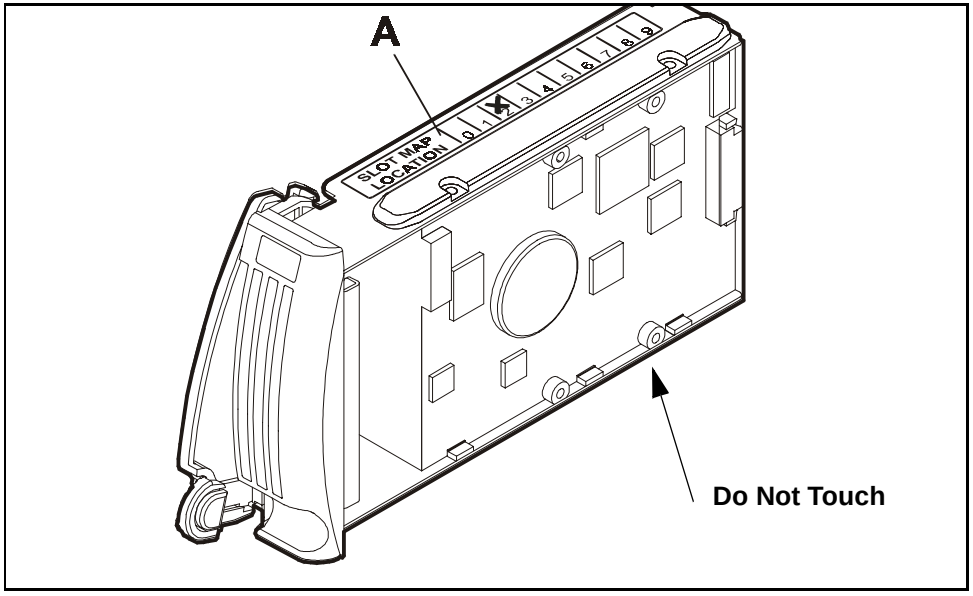



**Figure 61** Disk Module Slot Location Map Label

6. Open the cam latch (C Figure 60) by pulling the tab toward you.
7. Align the disk insertion guide (F) with a slot guide (G) and insert the disk into the slot. Typically, install disk modules on the left side of the enclosure and fillers on the right

**Note** Installing disks left to right allows you to insert the disk completely without releasing your grip on the handle.

8. Push the disk all the way into the chassis, letting the internal guides control the angle.
9. Lock the disk in place by pushing the latch toward the disk until it clicks.
10. Repeat steps 3–9 to install additional disks and fillers.

***CAUTION*** *Every slot must contain either a disk or filler.*

## Moving a Disk Enclosure from One Disk Array to Another

If you have more than one HP SureStore E Disk Array FC60, it is possible to move a disk enclosure from one array to another to balance capacity. The steps involved in moving a disk enclosure are the same for installing a new enclosure.

The following steps describe moving disk enclosures from one disk array to another.

1. If necessary, backup all data on the disks in the enclosures to be moved.
2. From the host, stop all I/O activity to the disks in the enclosures and unmount the file systems from the disks.
3. Unbind all LUNs on the disks in the enclosures using the management tool of your choice:
   - To use SAM, see "Unbinding a LUN" on page 271
   - To use Array Manager 60, see "Calculating LUN Capacity" on page 292
   - To use STM, see "Unbinding a LUN" on page 315.

**Note** **What if LUN 0 is on disks in the enclosure?**
If any of the disks in the enclosure are part of LUN 0, you will not be able to unbind the LUN before moving the disks. Instead, you must replace LUN 0 and exclude any of the disks in the enclosures from LUN 0.

4. Power down the disk array, both the array controller enclosure and the disk enclosures.
5. Remove the disk enclosures from the rack. To remove the enclosures, reverse the steps listed in "Installing the Disk Enclosures" on page 160.
6. Install the disk enclosures in the new rack following the instructions in "Installing the Disk Enclosures" on page 160. Make sure the disk enclosure is configured properly for the new disk array.
7. Bind a LUN with the new disks using the management tool of your choice:
   - To use SAM, see "Binding a LUN" on page 267
   - To use Array Manager 60, see "Binding a LUN" on page 289
   - To use STM, see "Binding a LUN" on page 314

**Note** After binding a LUN, you must execute the `insf -e` command to install special device files on the LUN. This makes the LUN usable by the operating system.

# Installing the Controller

This procedure describes how to install the Disk Array FC60 controller enclosure into an HP legacy rack or an HP System/E Rack.

## Step 1: Gather Required Tools

- Torx T25 screwdriver
- Torx T15 screwdriver
- Small flat-blade screwdriver

## Step 2: Unpack the Product

1. Lift off the overcarton and verify the contents (see Table 23 and Figure 62).

**Table 23** Controller Package Contents

| Figure Label | Part Description (See) |
|---|---|
| A | Controller chassis with pre-installed modules |
| B | Filler panel, 1/2 EIA unit, 2ea. |
| C | Rail kit (A5251A) for System/E racks |
| D | Rail kit (A5250A) for legacy racks, 1 ea. |
| E | SCSI Cables (length depends on option ordered)<br>2 meter (5064-2492) or 5 meter (5064-2470)<br>2 ea. / 1 disk enclosure 4 ea. / 4 disk enclosures<br>4 ea. / 2 disk enclosures 5 ea. / 5 disk enclosures<br>6 ea. / 3 disk enclosures 6 ea. / 6 disk enclosures |
| F | LVD SCSI Terminators (# of terminators / # disk enclosures ordered)<br>4 ea. / 1 disk enclosure 2 ea. / 4 disk enclosures<br>2 ea. / 2 disk enclosures 1 ea. / 5 disk enclosures<br>0 ea. / 3 disk enclosures 0 ea. / 6 disk enclosures |
| G | Fibre Channel cables, 2 ea.; if ordered<br>2M (1005-0585)<br>16M (1005-0586) |
| H | Media Interface Adapter (MIA) 1 ea./ controller module (5064-2464) |
| I | Power cords, ferrite bead (5064-2482), 2 ea. |
| J | ESD strap (9300-2170) 1 ea. |
| K | User's manual & Quick Install Guide |

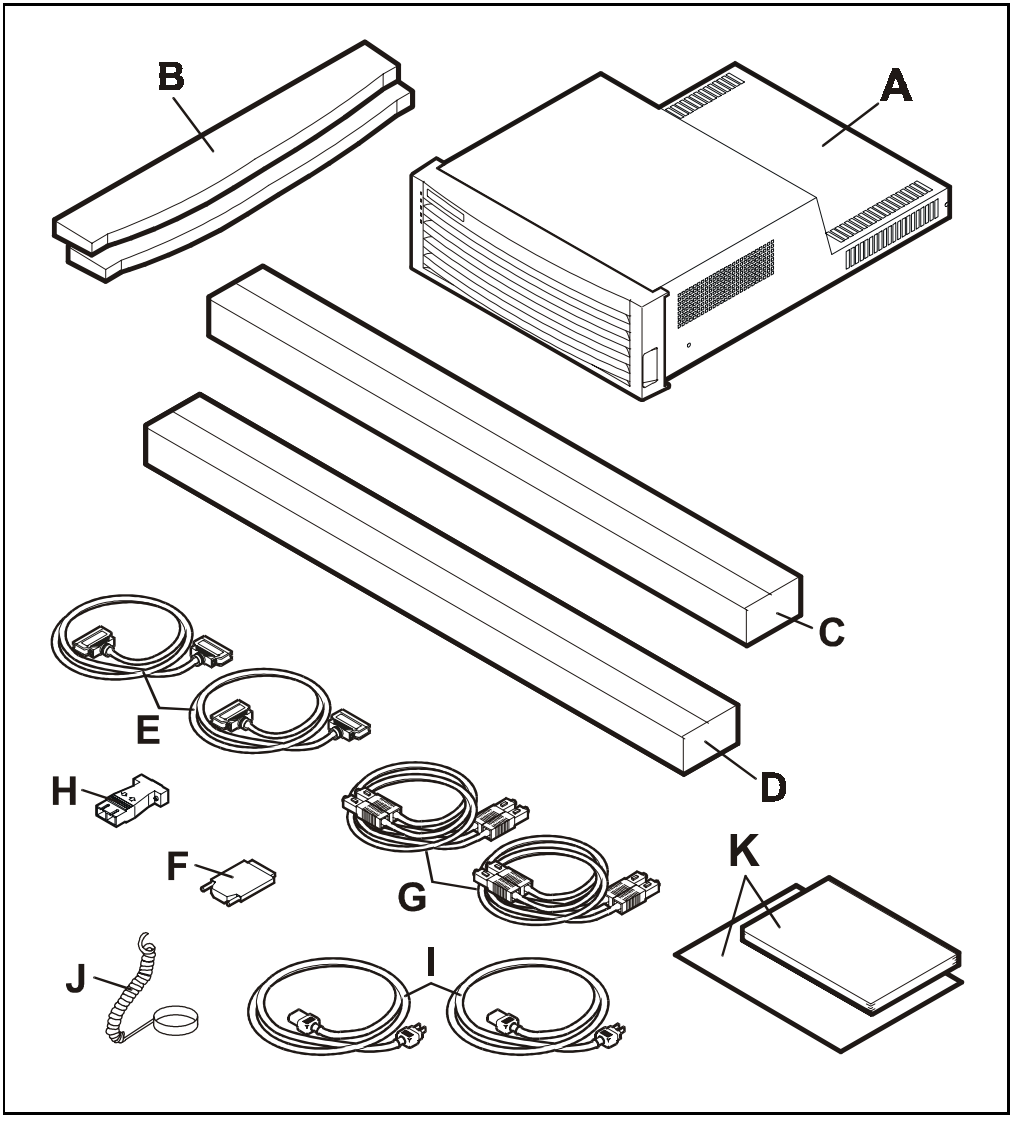



**Figure 62** Controller Enclosure Package Contents

## Step 3: Install Mounting Rails

Select the rail kit for the appropriate rack and follow the instructions included with the rail kit to install the rails in the rack. The following rail kits are available for use with the controller enclosure:

- HP A5250A for legacy HP Racks ((C2785A, C2786A, C2787A A1896A, or A1897A)
- HP A5251A for HP Rack System/E
- HP 5656A for Rittal 9000 racks

## Step 4: Install the Controller Enclosure

**Note** Prior to installing the controller enclosure, it may be easier to preform some of the installation steps described later in this chapter. For example, you may want to connect the MIA's, perform the Fibre Channel Host ID switch settings, and connect the SCSI terminators.

1. Remove the front cover by pulling out on the bottom and sliding it downward, slightly, then pulling it away from the enclosure.

**WARNING** **Do not attempt to lift the controller enclosure without the help of another person or a lift device. It weighs 75 pounds.**

2. With another person or using a lift, lift the enclosure and slide it, back-end first, into the rack, onto the previously installed rails (Figure 63). Push the enclosure as far into the rack as it will go.
3. Verify that two holes in each mounting ear (B in Figure 63) align with the sheet metal nuts previously installed on the rack front columns.
4. Insert two screws (A in Figure 63) through the holes in the controller mounting bracket that align with the nuts in rack column, on each side of the enclosure. Tighten the screws.

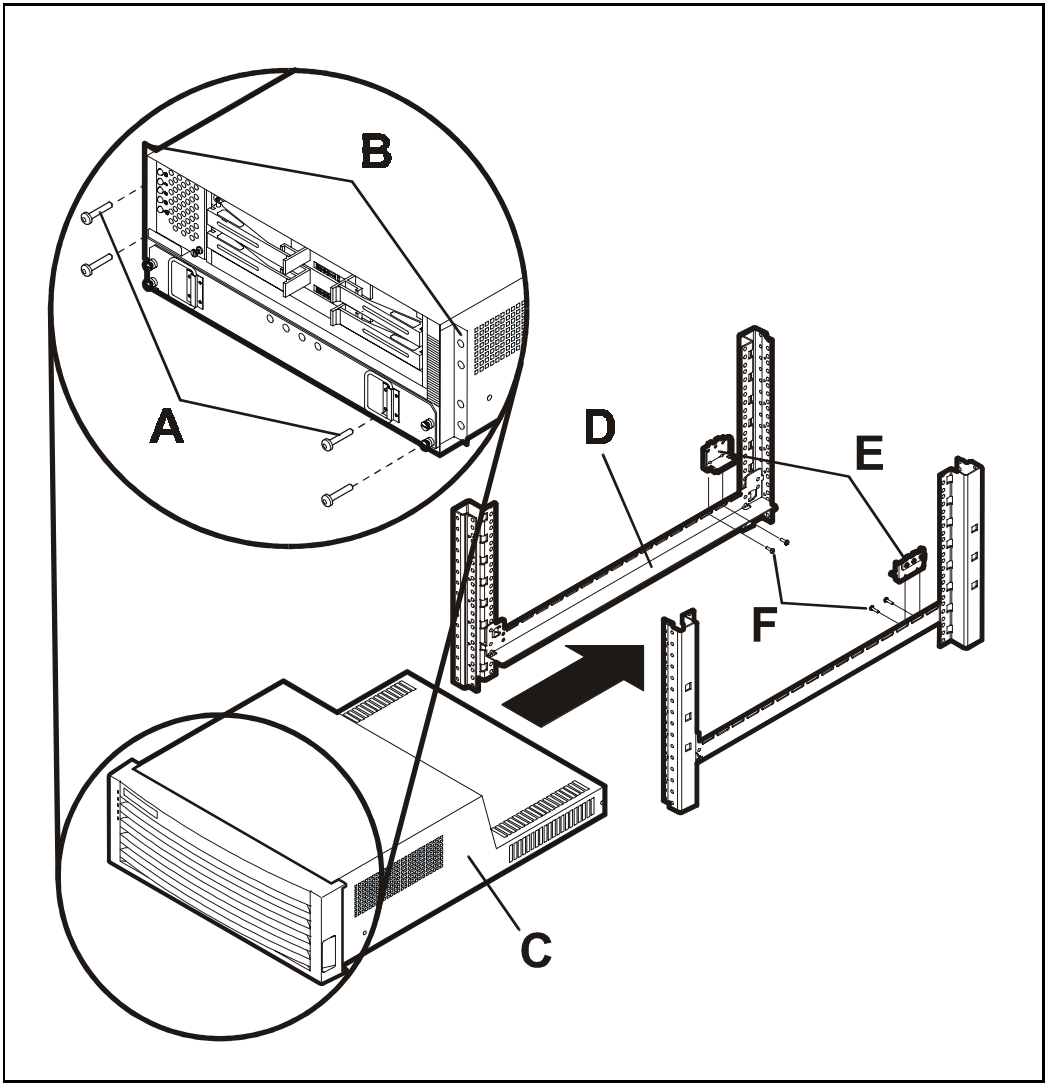



**Figure 63** Mounting the Controller Enclosure

5. If installing in an HP rack, secure the back of the enclosure to the rails using the two rail clamps from the rail kit.

   **In legacy HP racks:**

   a. Align screw holes and insert the clamp tab into the slot in the upper surface of the rail.

   b. Insert a screw through the hole in the clamp and the rail and tighten with a Torx T25 screwdriver.

   **In HP Rack System/E racks:**

   a. Set the clamp (E in Figure 63) inside the rail with the holes in the clamp along the slots in the rail.

   b. Push the clamp tight against the back of the disk enclosure. The curved tip of the clamp should overlap the bottom edge of the power supply slot.

   c. Insert and tighten two M5 16mm screws (F in Figure 63) through each clamp and rail.

# Configuration Switches

This section describes the configuration switches on the controller enclosure and the disk enclosures. Configuration switch settings must be set to support the configuration (full-bus ore split-bus) being installed (as planned for from chapter 2, Topology and Array Planning). Controller enclosure and the disk enclosures configuration switches include:

- Disk enclosure (Tray) ID switches
- Disk enclosure DIP switches
- Controller enclosure Fibre Channel host ID address switches

### Tips for Setting Configuration Switches

- All switch settings should be the same on both BCC modules on the disk enclosures.
- Disk enclosures must be power cycled before the new switch setting will take effect.
- If a BCC is installed that has settings different than its peer, a buzzer will sound. The settings must be corrected to use the disk enclosure.

## Disk Enclosure (Tray) ID Switch

The disk enclosure ID switch sets the Enclosure ID (or Tray ID). This value is used by FC60 diagnostic software and the controller to identify the disk enclosure for status monitoring operations. Each enclosure in the array must be set to a different Enclosure ID. To set the ID, set the rotary ID switches on the back of both BCC modules to the same number. See Figure 64. Typically, the disk enclosure closest to the controller enclosure is set to “0”, the next one down to “1”, and so on down the rack. See "Selecting Disks for a RAID 0/1 LUN" on page 243 for additional disk module addressing information.)

## Disk Enclosure DIP Switch

A DIP switch located on the back of each BCC is used to set configuration options for BCC operation. See Figure 64. These DIP switches contain five switch segments. Switch segment 1 is used to select either split-bus or full-bus operation. See "Full-Bus/Split-Bus (Switch 1) Configuration" on page 178. The remaining four switch segments are not used for disk array operation and must always be set to “0.”

**Note** Note that one BCC is inverted with respect to the other. Thus, the settings on one BCC appear as inverted and in reverse order from the other.

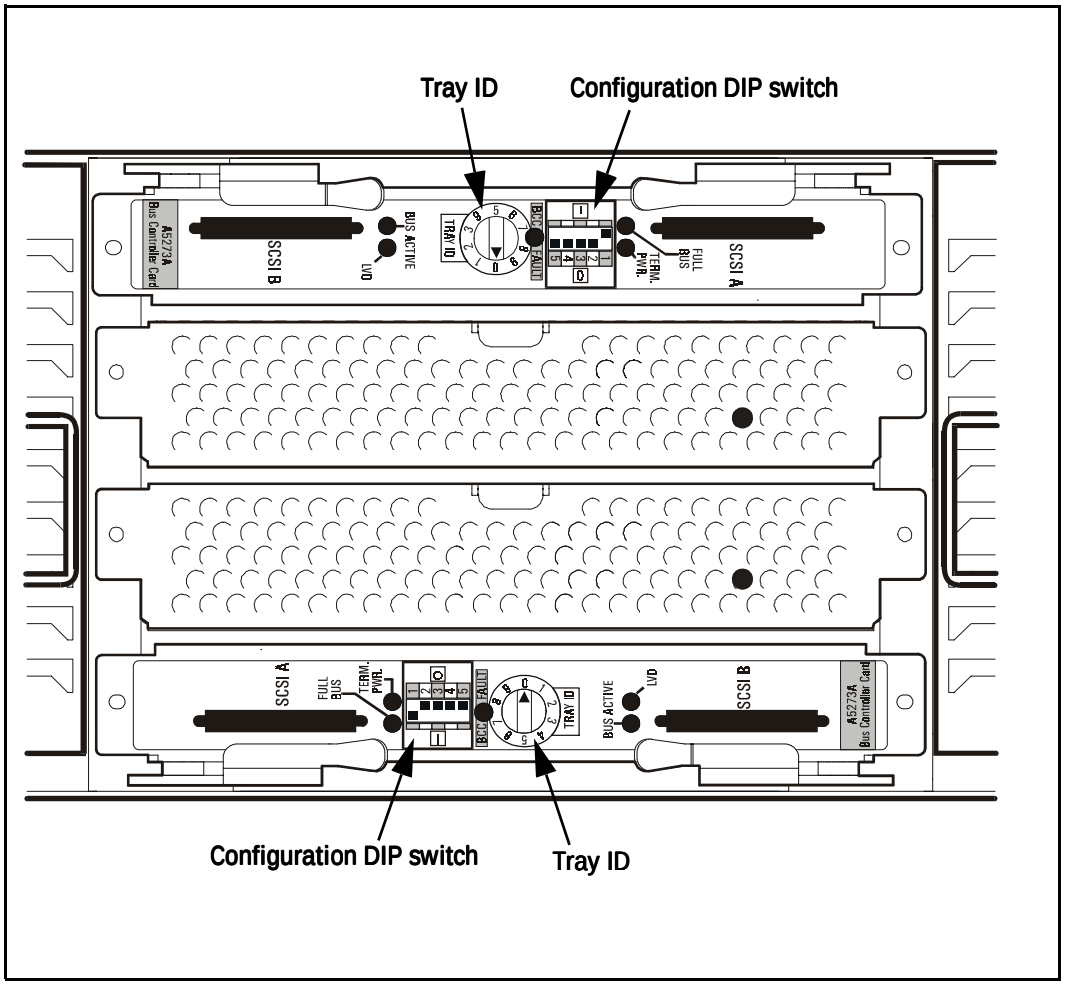



**Figure 64** Disk Enclosure Configuration Switches

**Table 24** Disk Enclosure Switches

| Switch | Setting | Operation |
|---|---|---|
| 1 - Full-/Split Bus Mode | 1 | Full-Bus Mode |
| | 0 | Split-Bus Mode |
| 2 - Stand-Alone/Array Mode | 0 | Always set to Off (Array Mode) |
| 3 - Bus Reset on Power Fail | 0 | Must be set to 0 |
| 4 - High/Low Bus Addressing | 0 | Set to 0 (Low addressing) |
| 5 - Not Used | 0 | Not used; must be set to 0 |

## Full-Bus/Split-Bus (Switch 1) Configuration

The disk enclosure's internal bus connects the disk drives together and to the BCCs. This bus can be set for two modes of operation: full-bus or split-bus mode. When switch 1 is set to Full Bus (1), all ten disk can be accessed together on a single SCSI bus. In this mode, one SCSI channel from the FC60 controller is connected to all ten disk modules through either BCC (typically BCC A).

**Note** In full bus mode, all ten disks can be accessed through either BCC. However, internally each BCC still manages five disks. This means that if the BCC that is not connected to a SCSI channel fails, access to its five disks will be lost. Failure of the BCC that is connected to a SCSI channel will render all ten disks inaccessible.

When switch 1 is set to 0 (split-bus mode), the disk enclosure is split into two separate buses with five disks to each bus. In split-bus mode, each bus is controlled by a separate BCC. BCC A (the top BCC) controls the even drive slots (0, 2, 4, 6, 8) and BCC B (the bottom BCC) controls the odd bus slots (1, 3, 5, 7, 9). In split-bus mode, two SCSI channels are connected to the disk enclosure, one to a SCSI connector on BCC A and a second a SCSI connector on BCC B.

Each slot in the disk enclosure is assigned a SCSI ID. A range of IDs is assigned to the odd disk slots, and a different range to the even slots (see Table 25). The two ID ranges include

a low range of IDs (0, 1, 2, 3, and 4) and a high range of IDs (8, 9, 10, 11, and 12). (BCCs are also provided addresses as shown in Table 25). Note that the SCSI IDs do not correspond to the physical slot number.

The assignment of the SCSI IDs differs depending on whether the enclosure is operating in full-bus or split-bus mode. When full-bus mode is selected, the low ID range (0 - 4) is assigned to the even disk slots, and the high range (8 - 12) is assigned to the odd slots. See Table 25. When the disk enclosure is in split-bus mode, the low ID range is assigned to both the even slots and the odd slots. This is possible because the two busses are isolated within the enclosure. See "Selecting Disks for a RAID 0/1 LUN" on page 243 for additional disk module addressing information.

**Table 25** Disk Slot, Full-Bus/Split-Bus Mode, SCSI IDs

| | Physical Disk Slot #s | | | | | | | | | | BCC | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|
| | 0 | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | A | B |
| **Full-Bus Mode SCSI IDs** | | | | | | | | | | | | |
| BCC A (or B) | 0 | 8 | 1 | 9 | 2 | 10 | 3 | 11 | 4 | 12 | 14 | 15 |
| **Split-Bus Mode SCSI IDs** | | | | | | | | | | | | |
| BCC-A (Even Slots) | 0 | - | 1 | - | 2 | - | 3 | - | 4 | - | 15 | - |
| BCC-B (Odd Slots) | - | 0 | - | 1 | - | 2 | - | 3 | - | 4 | - | 15 |

## Fibre Channel Host ID Address Setting

The controller enclosure has two Fibre Channel connections, one for each controller module. Each controller module communicates with the host through its Fibre Channel loop connection. Internally the two controllers are connected to a common Fibre Channel loop. Because they are connected to a common loop, each controller module must have a unique Fibre Channel Host ID.

The Fibre Channel Host IDs are selected by two separate DIP switches located on the back of the controller enclosure. See Figure 65. Switch Host ID BD1 SW1 selects the address for

controller module A (Fibre Channel connector J3) and Host ID BD2 SW2 selects the address for controller module B (Fibre Channel connector J4).

Each Fibre Channel Host ID DIP switch contains a bank of seven switches that select the address using a binary value, 000 0000 (0) through 111 1111 (126). To set an address, set the switches in the up position for “1” or down for “0” (refer to Table 26 for binary switch settings). Figure 65 illustrates the loop ID switch set to 42 (0101010).

**Note** The factory default Fibre Channel Host ID loop address settings for controller A (BD1) is 5, and the default for controller B (BD2) is 4. If you must change the default values, make sure you select a unique value for each ID. The Fibre Channel Host ID loop address settings must be unique for each controller, regardless of the external Fibre Channel topology.

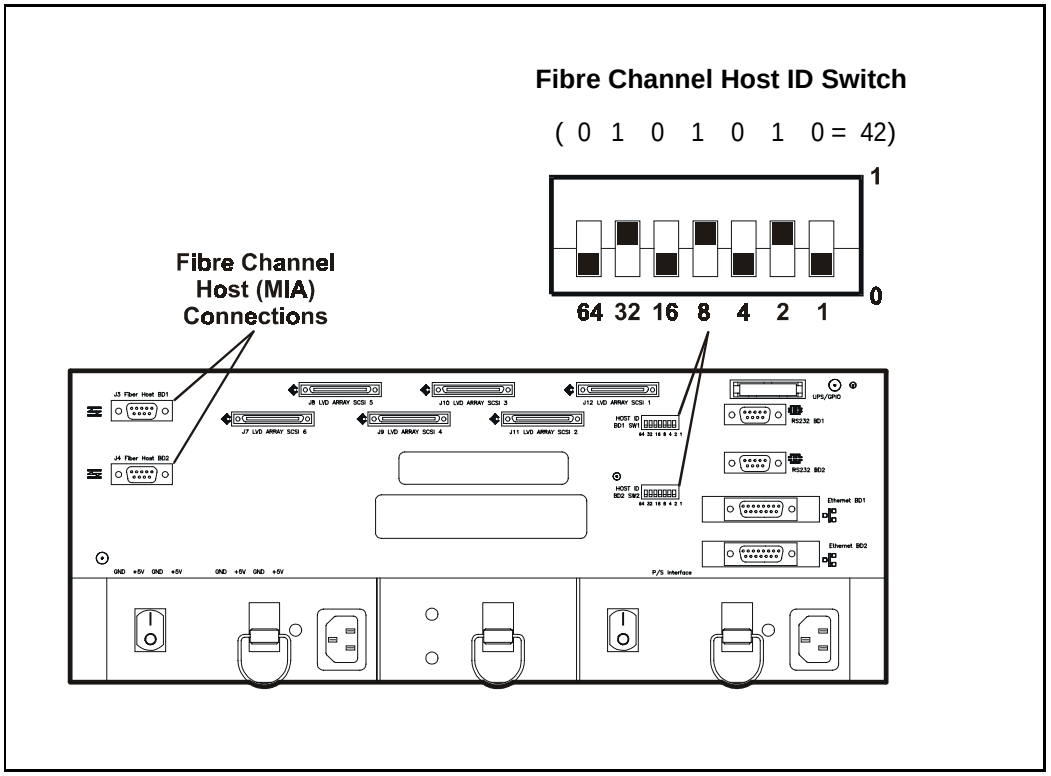



**Figure 65** Fibre Channel Connectors and Fibre Channel Host (Loop) ID Switches

**Note** Occasionally two or more ports in an arbitrated loop will arbitrate simultaneously. Priorities are decided according to the loop IDs. The higher the loop ID, the higher the priority.

**Table 26** Fibre Channel Addresses

| Decimal | Binary | Decimal | Binary | Decimal | Binary | Decimal | Binary |
|---|---|---|---|---|---|---|---|
| 0 | 000 0000 | 38 | 010 0110 | 76 | 100 1100 | 114 | 111 0010 |
| 1 | 000 0001 | 39 | 010 0111 | 77 | 100 1101 | 115 | 111 0011 |
| 2 | 000 0010 | 40 | 010 1000 | 78 | 100 1110 | 116 | 111 0100 |
| 3 | 000 0011 | 41 | 010 1001 | 79 | 100 1111 | 117 | 111 0101 |
| 4 | 000 0100 | 42 | 010 1010 | 80 | 101 0000 | 118 | 111 0110 |
| 5 | 000 0101 | 43 | 010 1011 | 81 | 101 0001 | 119 | 111 0111 |
| 6 | 000 0110 | 44 | 010 1100 | 82 | 101 0010 | 120 | 111 1000 |
| 7 | 000 0111 | 45 | 010 1101 | 83 | 101 0011 | 121 | 111 1001 |
| 8 | 000 1000 | 46 | 010 1110 | 84 | 101 0100 | 122 | 111 1010 |
| 9 | 000 1001 | 47 | 010 1111 | 85 | 101 0101 | 123 | 111 1011 |
| 10 | 000 1010 | 48 | 011 0000 | 86 | 101 0110 | 124 | 111 1100 |
| 11 | 000 1011 | 49 | 011 0001 | 87 | 101 0111 | 125 | 111 1101 |
| 12 | 000 1100 | 50 | 011 0010 | 88 | 101 1000 | 126 | 111 1110 |
| 13 | 000 1101 | 51 | 011 0011 | 89 | 101 1001 | 127 | 111 1111 |
| 14 | 000 1110 | 52 | 011 0100 | 90 | 101 1010 | | |
| 15 | 000 1111 | 53 | 011 0101 | 91 | 101 1011 | | |
| 16 | 001 0000 | 54 | 011 0110 | 92 | 101 1100 | | |
| 17 | 001 0001 | 55 | 011 0111 | 93 | 101 1101 | | |
| 18 | 001 0010 | 56 | 011 1000 | 94 | 101 1110 | | |
| 19 | 001 0011 | 57 | 011 1001 | 95 | 101 1111 | | |
| 20 | 001 0100 | 58 | 011 1010 | 96 | 110 0000 | | |
| 21 | 001 0101 | 59 | 011 1011 | 97 | 110 0001 | | |
| 22 | 001 0110 | 60 | 011 1100 | 98 | 110 0010 | | |
| 23 | 001 0111 | 61 | 011 1101 | 99 | 110 0011 | | |
| 24 | 001 1000 | 62 | 011 1110 | 100 | 110 0100 | | |
| 25 | 001 1001 | 63 | 011 1111 | 101 | 110 0101 | | |
| 26 | 001 1010 | 64 | 100 0000 | 102 | 110 0110 | | |
| 27 | 001 1011 | 65 | 100 0001 | 103 | 110 0111 | | |
| 28 | 001 1100 | 66 | 100 0010 | 104 | 110 1000 | | |
| 29 | 001 1101 | 67 | 100 0011 | 105 | 110 1001 | | |
| 30 | 001 1110 | 68 | 100 0100 | 106 | 110 1010 | | |
| 31 | 001 1111 | 69 | 100 0101 | 107 | 110 1011 | | |
| 32 | 010 0000 | 70 | 100 0110 | 108 | 110 1100 | | |
| 33 | 010 0001 | 71 | 100 0111 | 109 | 110 1101 | | |
| 34 | 010 0010 | 72 | 100 1000 | 110 | 110 1110 | | |
| 35 | 010 0011 | 73 | 100 1001 | 111 | 110 1111 | | |
| 36 | 010 0100 | 74 | 100 1010 | 112 | 111 0000 | | |
| 37 | 010 0101 | 75 | 100 1011 | 113 | 111 0001 | | |

# Attaching Power Cords

Each enclosure (controller and disk enclosures) contains dual power supplies that must be connected to the power source (PDU). When connecting power cords for high availability installations, connect one enclosure power cord to one source (PDU) and the other power cord to an alternate source (PDU). To complete the power connection, follow the steps below.

1. Set the power switch on the disk enclosures to OFF. The switch is located at the front, top, right corner of the enclosure.
2. Set the two power switches on the back of the controller enclosure to OFF (0).
3. Connect power cords to both power supplies for each enclosure. Attach one end of the power cord to the enclosure power supply receptacle and the other end to the PDU/ PDRU receptacle.

**Note** The controller enclosure requires special ferrite bead power cords (part no. 5064-2482) for FCC compliance.

   When connecting power cords, connect the power cords from the right-side power supplies to the PDUs/PDRUs on the right side of the rack and connect the power cords from the left-side power supplies to the PDUs/PDRUs on the left side of the rack. This method ensures that each enclosure power supply is connected to a separate power source to ensure high availability.

4. Connect the PDUs/PDRUs into their respective power sources.

   Choose outlets according to the following guidelines:

   – **Redundancy**. To extend the redundancy of the product, attach each cord to a different PDU. This is represented in Figure 66 and Figure 67 by the absence of duplicate letters in each disk enclosure.

   – **Reliability** To avoid cascading faults for a group of disk enclosures that are plugged into the same PDU, distribute the redundant power cords among different PDUs. This is represented in Figure 66 and Figure 67 by the absence of duplicate pairs of

letters among all disk enclosures. “Cascading” refers to overload faults that occur on a backup PDU as a result of power surges after the primary PDU fails.

– **Serviceability**. Choose PDU locations that prevent power cords from interfering with the removal and replacement of serviceable components. Also leave a 6-inch service loop to allow for the rotation of PDRUs.

The letters A, B, C, D, E and F in Figure 66 and Figure 67 represent independent PDUs or PDU banks. The absence of duplicate letters in individual sorage systems indicates the products are using redundant PDUs. The minimum number of duplicate pairs of letters in all sorage systems indicates the products are protected against cascading faults.

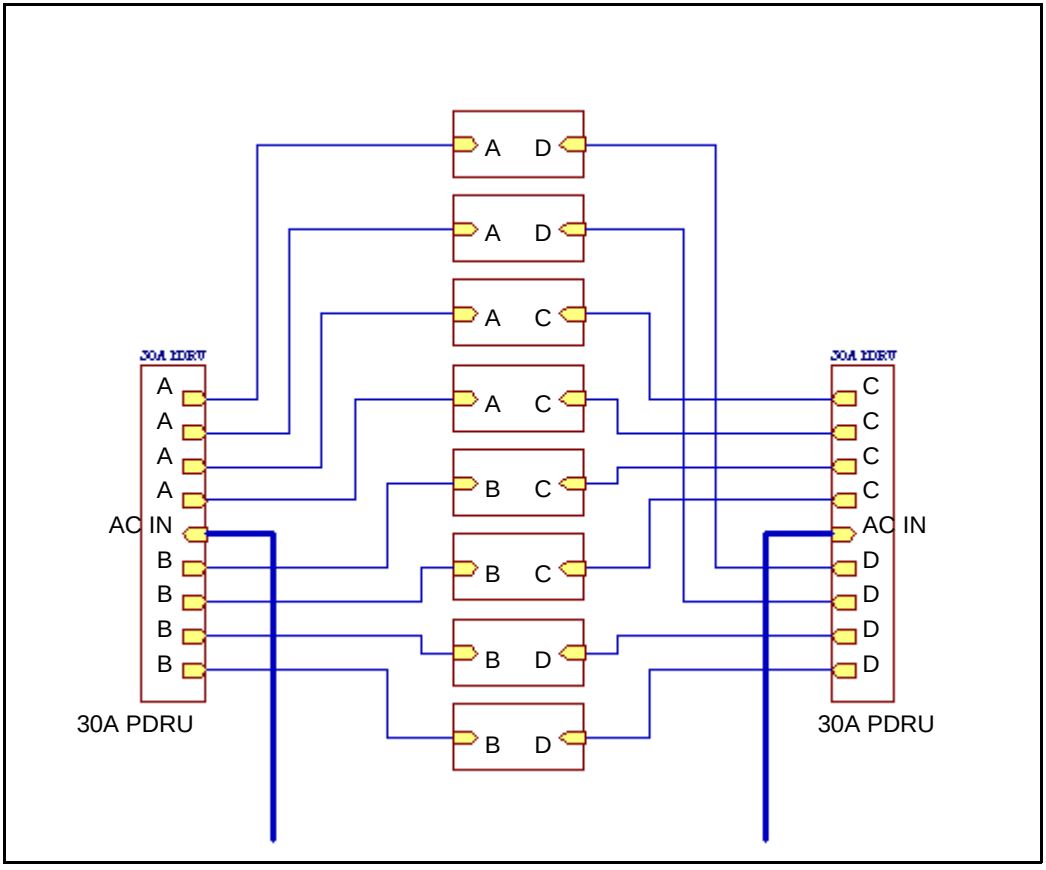



**Figure 66** Wiring Scheme for 1.6-Meter Rack

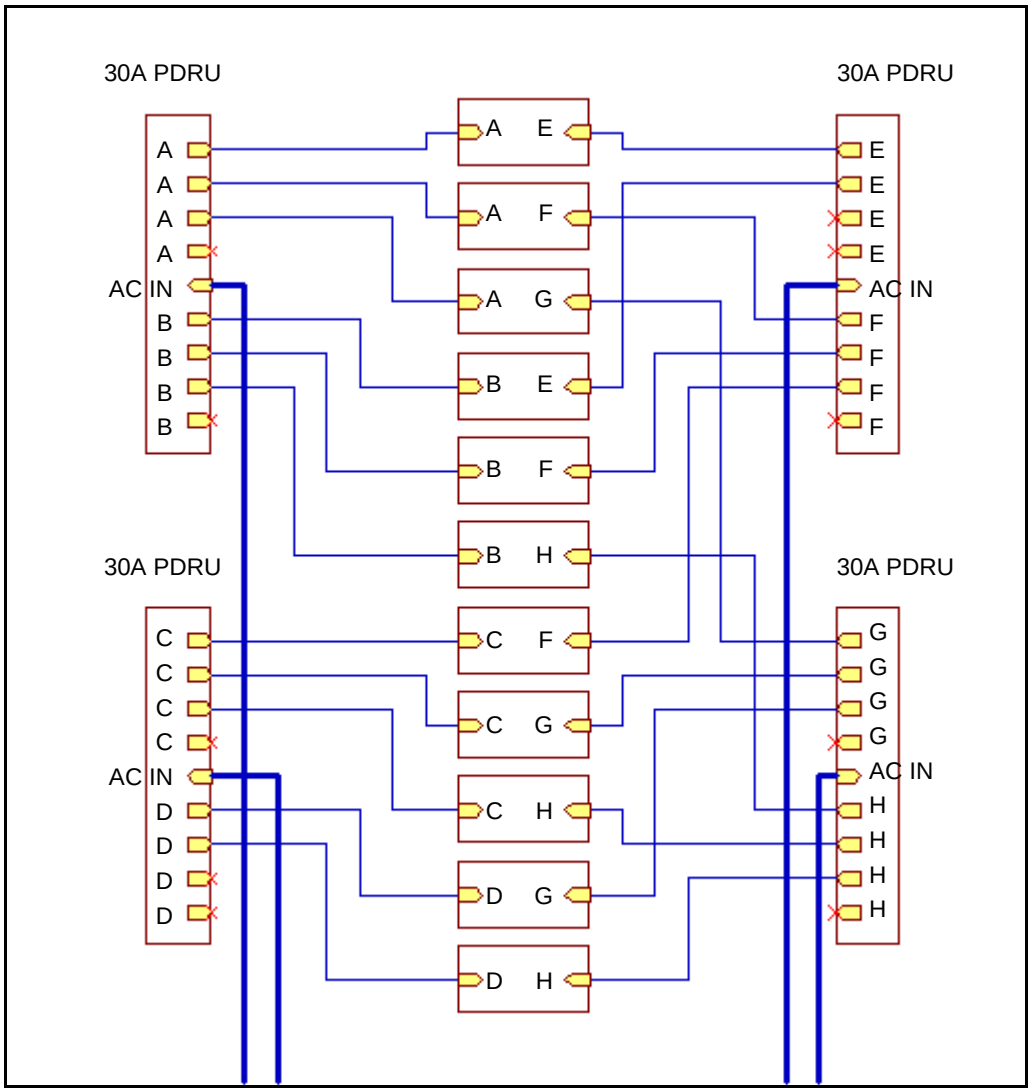



**Figure 67** Wiring Scheme for 2.0-Meter Rack

# Attaching SCSI Cables and Configuring the Disk Enclosure Switches



**NOTE!** It is critical that all SCSI cables be tightened securely. Use the following steps to ensure the cable connectors are seated properly.

1. Connect the cable to the enclosure connector and tighten the mounting screws finger tight.
2. Push on the connector and retighten the mounting screws. Repeat once more.
3. Use a flat blade screwdriver to tighten the screw appropriately. Be sure the screw is not cross-threaded.



There are two types of cabling configurations for the array: full-bus and split-bus. The configuration you select is determined by a number of factors which are described in Chapter 2, Topology and Array Planning.

When you have determined which configuration you will use, refer to "Full-Bus Cabling and Switch Configuration" on page 188" or "Split-Bus Switch and Cabling Configurations" on page 191".

## A Word About SCSI Terminators

Terminators need to be installed on both the disk enclosures and the controller enclosure. On the controller enclosure, any unused SCSI connector must be terminated. On the disk enclosure, not every connector requires termination. If a SCSI cable is connected to a BCC connector, the other connector on the BCC must be terminated. If there is no SCSI cable connected to the BCC, no terminator is required on either connector. A BCC that is operating in split-bus mode will have a cable connected to each BCC. In this case, the unused connector on each BCC must be terminated. See Figure 69 or Figure 71.

## Full-Bus Cabling and Switch Configuration

Cabling for a full bus configuration requires connecting one SCSI cable from the controller to the disk enclosure and setting the configuration switches. Figure 69 illustrates full-bus cabling connections for a six disk enclosure array. It is possible to configure any number of disk enclosures, from one to six, using this method. However, full bus is typically used when four or more disk enclosures are installed in the array.

**Note** For consistency and ease of managment, it is recommended that you observe the following rules:

- Connect channel 1 on the controller enclosure to the disk enclosure mounted directly beneath the controller enclosure. Connect channel 2 to the next disk enclosure down, and so on down the rack.

- Install cables along the left side of the disk enclosure connectors (as viewed from the back of the rack) and install terminators along the right side connectors.

To configure the disk array for full-bus operation:

1. Connect one end of a SCSI cable to the controller enclosure SCSI channel 1 connector, J12 JVD ARRAY SCSI 1. See Figure 69.
2. Connect the other end of the SCSI cable to the disk enclosure BCC A, SCSI connector B.

**Note** The lower BCC is inverted from the upper BCC. Make sure you take this into account when setting the switches.

3. Set the Enclosure (Tray) ID rotary switches on both BCCs to “0”. See Figure 68.
4. Set the Enclosure DIP switches on both BCCs to the following settings:
   a. Set segment 1 to “1”
   b. Set the remaining four segments to “0”. See Figure 68.
5. Connect a terminator to SCSI connector A on BCC A (see Figure 69, “**T**”).
6. Repeat the preceding steps for each disk enclosure, using the next sequential SCSI channel and incrementing the Disk Enclosure (Tray) ID setting for each enclosure.

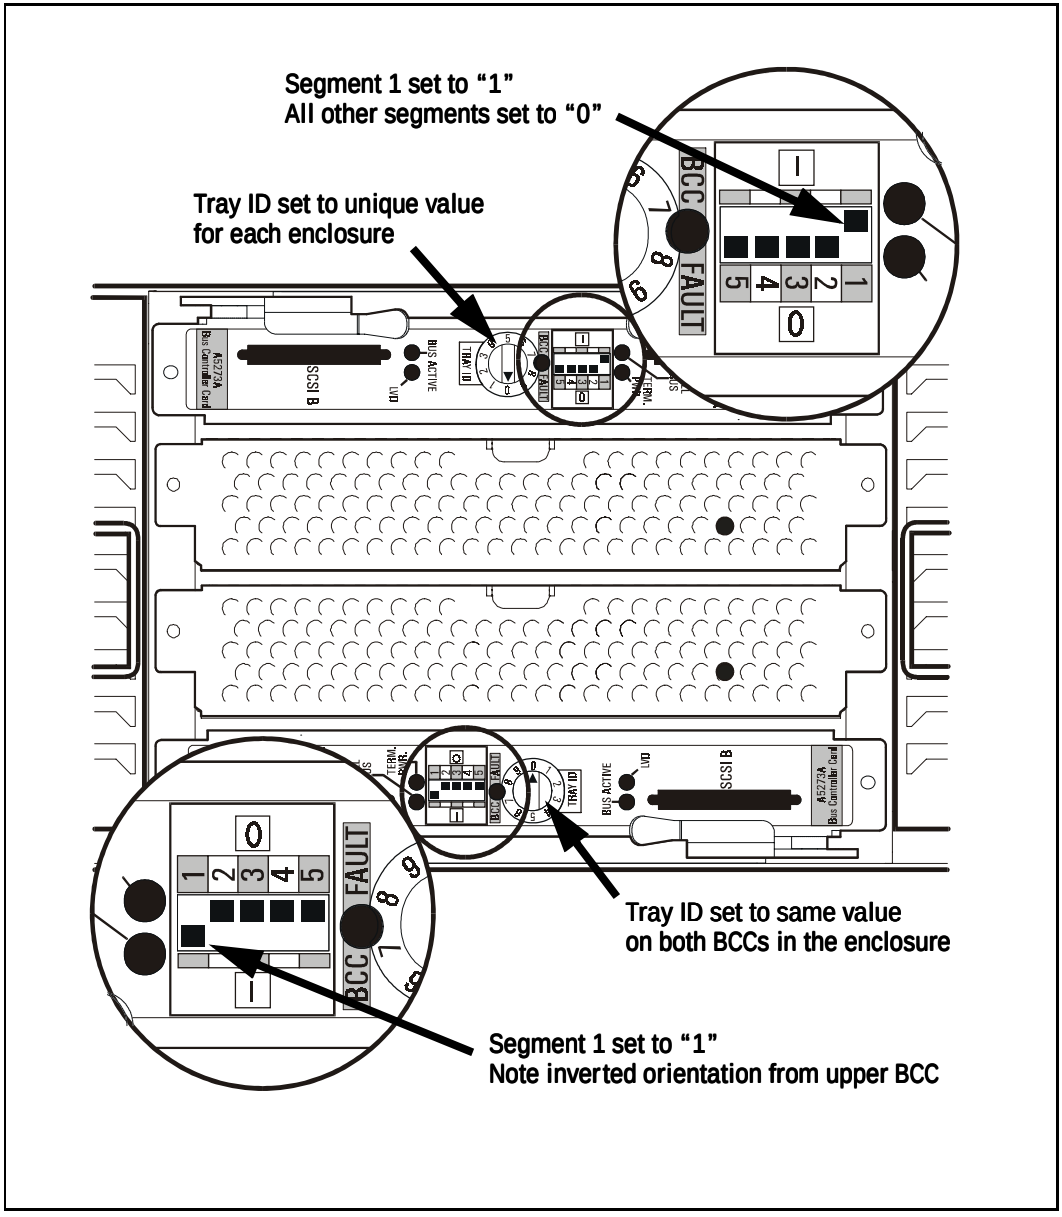

Segment 1 set to “1”
All other segments set to “0”
Tray ID set to unique value
for each enclosure
BCC FAULT
SCSI B
BUS ACTIVE
LVD
TRAY ID
Tray ID set to same value
on both BCCs in the enclosure
Segment 1 set to “1”
Note inverted orientation from upper BCC

**Figure 68** Full Bus BCC Configuration Switch Settings

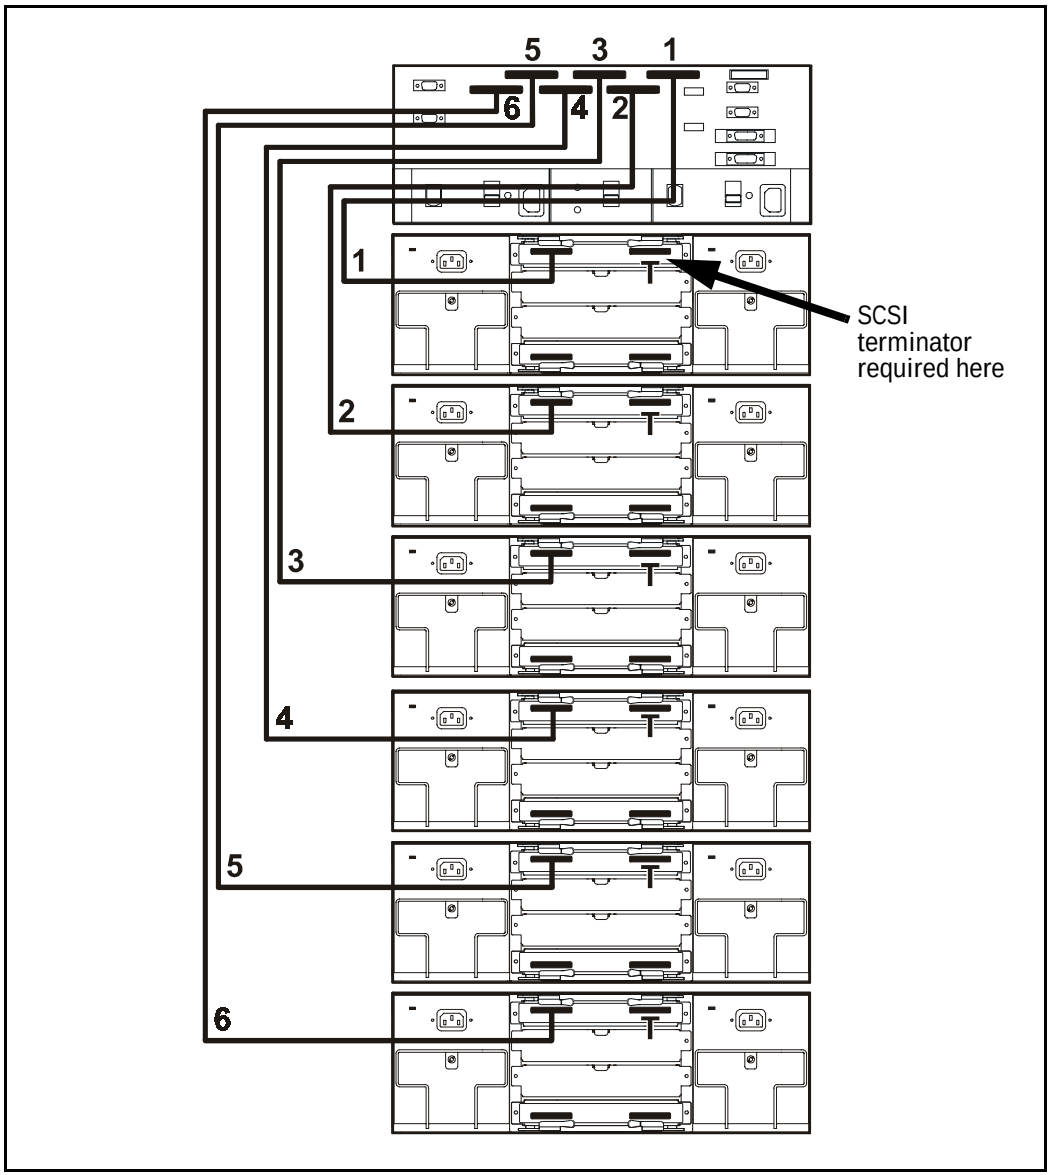



**Figure 69** Full-Bus Cabling

## Split-Bus Switch and Cabling Configurations

Split-bus cabling requires two SCSI cables from each disk enclosure to the controller enclosure. Split-bus cabling is typically used for installations with 3 or fewer disk enclosures. Cabling for a split-bus configuration is shown in Figure 71.

**Note** For consistency and ease of managment, it is recommended that you observe the following rules:

- Connect channels 1 and 2 on the controller enclosure to the disk enclosure mounted directly beneath the controller enclosure. Connect channels 3 and 4 to the next disk enclosure down, and channels 5 and 6 to the next enclosure down.

- Install cables along the left side of the disk enclosure connectors (as viewed from the back of the rack) and install terminators along the right side connectors.

To configure the disk array for split-bus operation:

1. Connect one end of a SCSI cable to the controller enclosure SCSI channel 1 connector, J12 JVD ARRAY SCSI 1 (see Figure 71).
2. Connect the other end of the SCSI cable to the disk enclosure BCC A, SCSI connector B.
3. Connect a second SCSI cable to the controller enclosure SCSI channel 2 connector, J11 JVD ARRAY SCSI 2.
4. Connect the other end of the SCSI cable to the disk enclosure BCC B, SCSI connector A .

**Note** The lower BCC is inverted from the upper BCC. Make sure you take this into account when setting the switches.

5. Set the Enclosure (Tray) ID rotary switches on both BCCs to “0”. See Figure 70.
6. Set all segments of the Enclosure DIP switches on both BCCs to “0”.
7. Connect a SCSI terminator to the remaining SCSI connector on each BCC.
8. Repeat the preceding steps for the remaining disk enclosures, using the next sequential SCSI channels and incrementing the Disk Enclosure (Tray) ID setting for each enclosure.

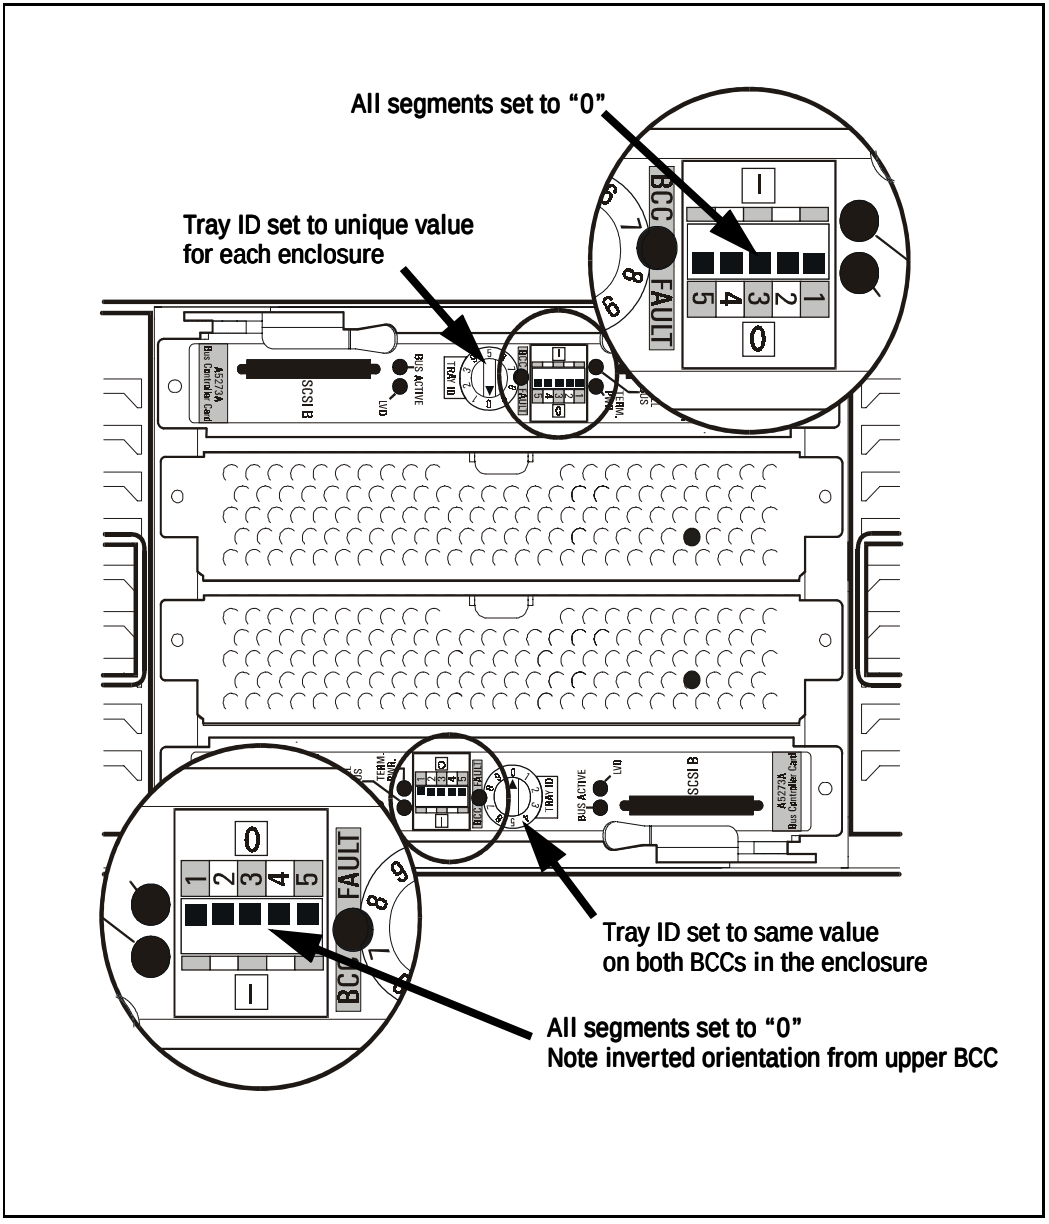



**Figure 70** Split- Bus Configuration Switch Settings

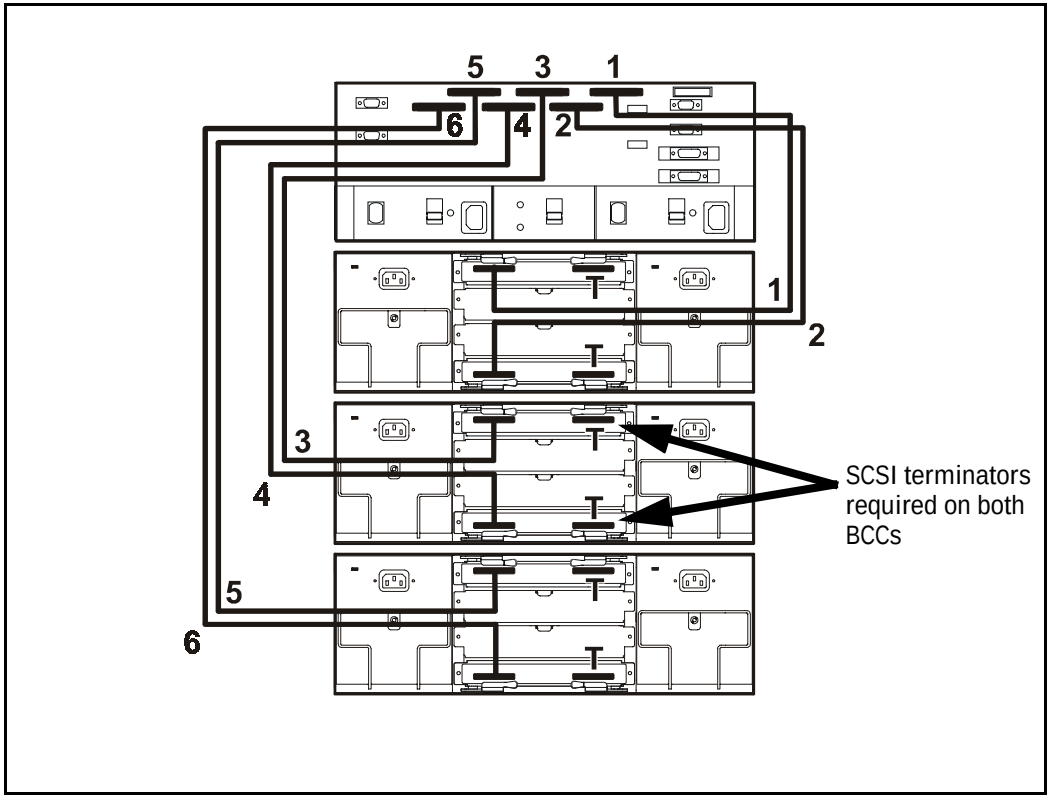



**Figure 71** Split-Bus Cabling

## Bus Addressing Examples

Each disk module within the disk array is identified by its channel number and SCSI ID. These values will differ depending on which type of bus configuration is used for the disk enclosures. See "How are disk modules in the array identified?" on page 244 for more information. Figure 72 is an example of split-bus addressing. Figure 73 is an example of full-bus addressing.

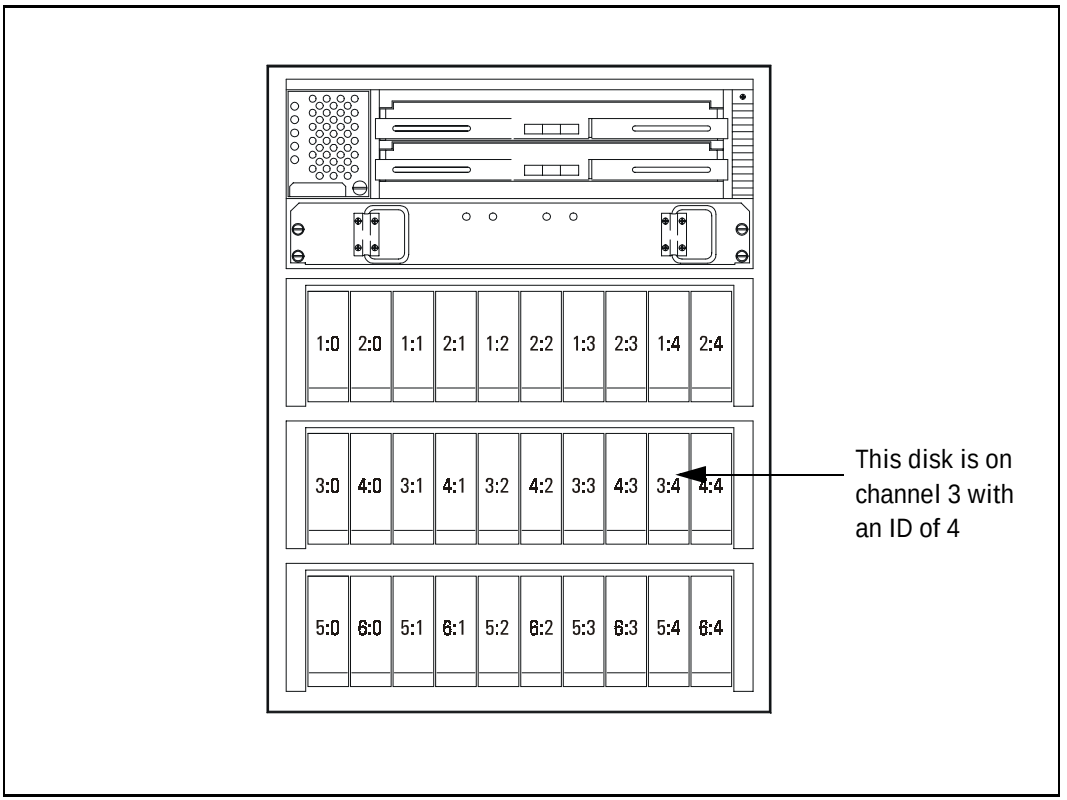



**Figure 72** Split-Bus Addressing Example

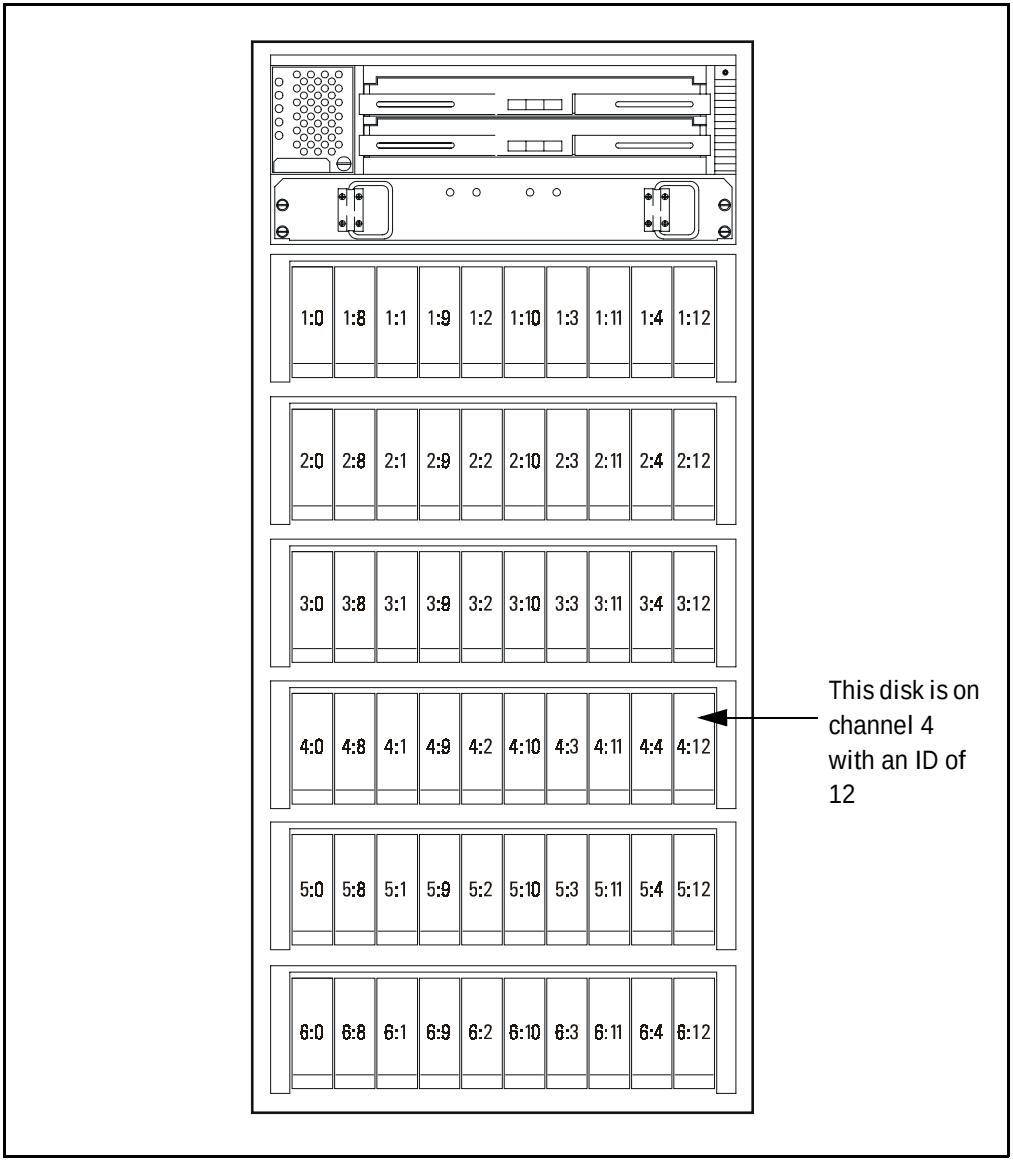



**Figure 73** Full-Bus Addressing Example

# Connecting the Fibre Channel Cables

Fibre Channel cables provide the I/O path to the disk array. The Fibre Channel cable connects the controller enclosure directly to a host, or to a hub.

For operation on HP-UX, the host must contain an HP Fibre Channel Mass Storage/9000 adapter. For HP-UX supported adapters, installation, and configuration information for the HP Fibre Channel adapters, refer to the *Hewlett-Packard Fibre Channel Mass Storage Adapter Service and User Manual* (J3636-90002) supplied with the adapter, or you can download this document from http://www.docs.hp.com/. This document also contains information on verifying the Fibre Channel connection to the disk array and information on troubleshooting the Fibre Channel loop.

For information on Fibre Channel host adapters supported on Windows NT and Windows 2000, check the Host Adapter folder on the *HP Storage Manager 60 CD.*

To connect the Fibre Channel cables to the array, complete the following steps:

**Note** It is advisable to select Fibre Channel cables that are as close as possible to the proper length for the installation.

1. Connect a fibre optic cable to a host or hub. Locate the appropriate connector on the adapter card. Remove the optical cover and connect the fibre optic cable. Refer to the appropriate system documentation for additional information.

**Note** When plugging the Fibre Cable connector into the adapter, push the plug into the connector until it snaps into the connector. There is a slight amount of slack in the connection which may prevent the connector from making an optimal connection. After plugging the plug into the connector, grab the cable just behind the plug and push lightly inward to ensure a good optical connection.

2. Plug a media interface adapter (MIA) into the Fibre Channel connectors on the controller enclosure (Figure 74, connectors J3 and J4). Make sure that the RFI gaskets are installed between the MIA and the controller connector. Tighten the two screws in each MIA.

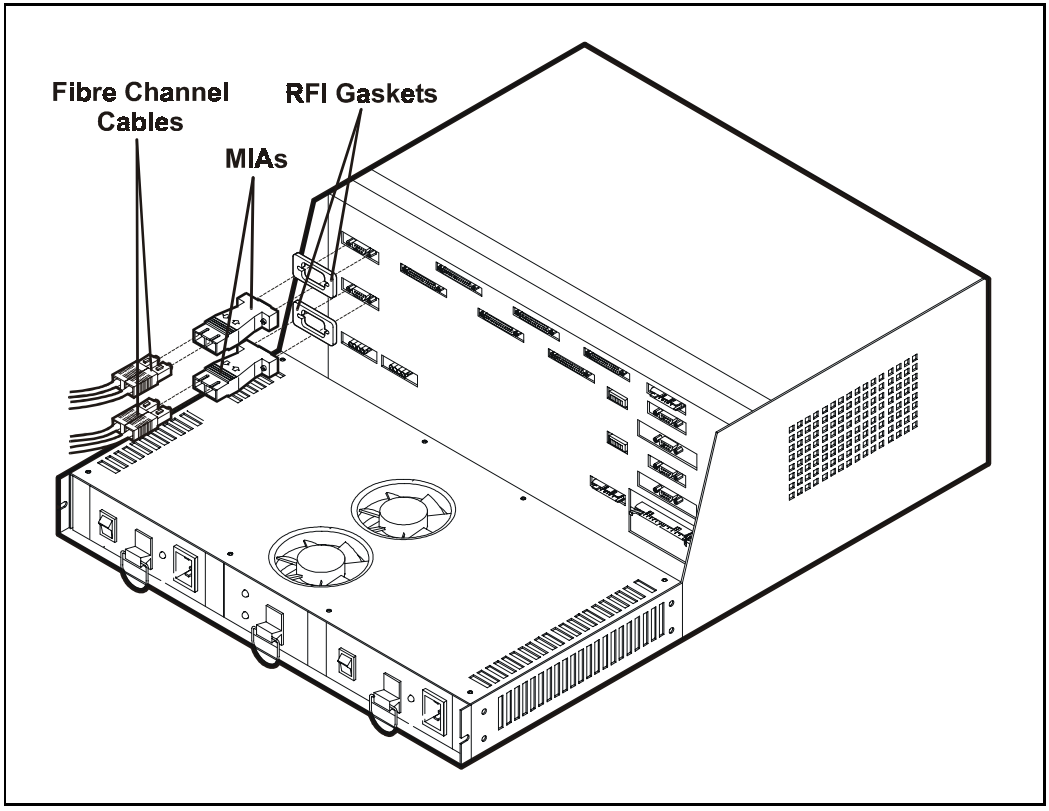



**Figure 74** MIA, RFI Gasket, and Fibre Channel Installation

3. Connect the Fibre Channel connectors to the controller MIAs.

   a. Remove the optical protectors from the ends of the MIAs and the Fibre Channel cables (Figure 74).

   b. Insert the Fibre Channel connectors into the MIAs. The fibre optic cable connector is keyed to install only one way.

# Applying Power to the Disk Array

Once the hardware installation is complete, the disk array can be powered up. It is important that the proper sequence be followed when powering up the components of the disk array. To ensure proper operation, power should be applied to the disk enclosures first and then to the controller enclosure, or all components can be powered up simultaneously. This gives the disks time to spin up, ensuring that the disks are detected when the controller comes on line.

Once power is applied to all components, proper operation of the disk array is verified by checking the LEDs on both the controller enclosure and disk enclosure.

**Note** **What happens if the disk enclosures are powered up before the controller enclosure?**
When the disk array controller enclosure is powered up, it checks for the presence of the disks in the disk enclosures. If the disk enclosures have not yet been powered on, the controller will not detect the disks. This will result in no LUNs appearing because none of the disks are visible to the controller. In this case it will be necessary to power down the disk array and power it up again in the proper sequence.

To power up the disk array:

**Note** The following procedure assumes the disk array components will be powered up individually. All disk array components can also be powered up simultaneously using a single rack power switch.

1. Open all front doors on the disk enclosures.
2. Power up each disk enclosure using the disk enclosure power switch (A in Figure 75).

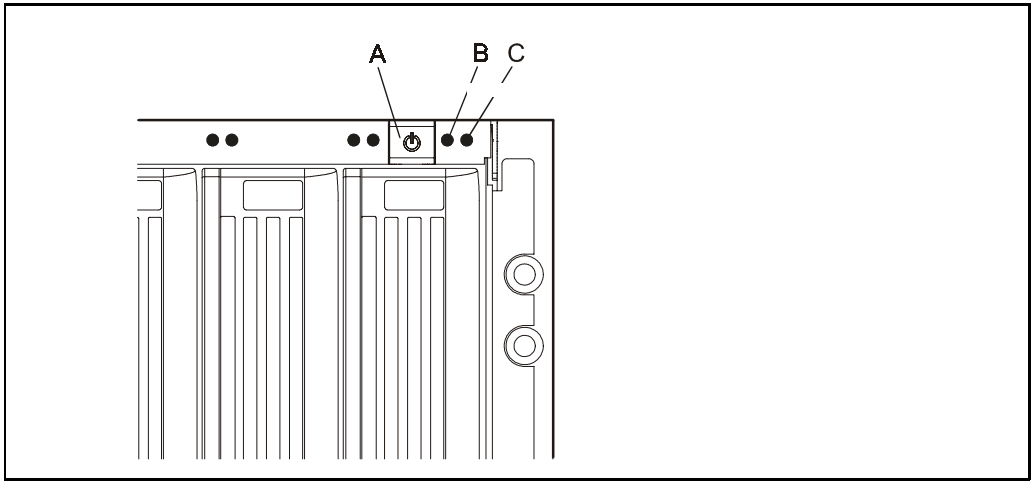



**Figure 75** Disk Enclosure Power Switch and System LEDs

3. Check the LEDs on the front of the disk enclosures (see Figure 77). The System Power LED (B in Figure 75) should be on and the Enclosure Fault LED (C) should be off.

   It is normal for the Enclosure Fault LED (amber) to go on momentarily when the enclosure is first powered on. However, if the Enclosure Fault LED remains on, it indicates that a fault has been detected. Refer to "Troubleshooting" on page 359 for additional information.

4. Power up the controller enclosure using both power switches on the back of the enclosure (see Figure 76).

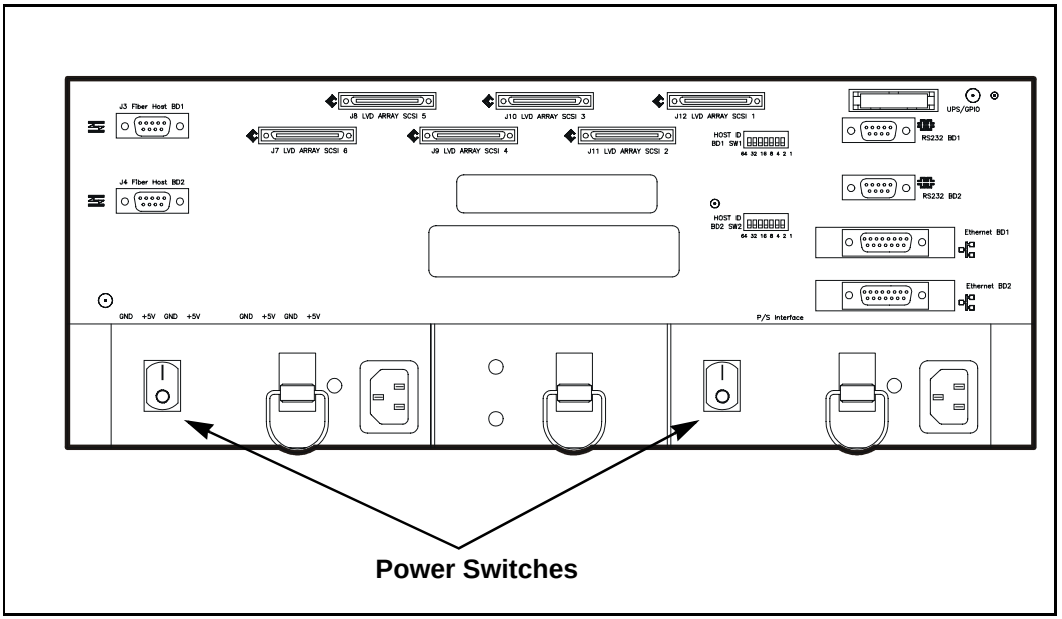



**Figure 76** Controller Enclosure Power Switches

5. Check the controller enclosure LEDs (see Figure 78). The Power LED should be on and the remaining LEDs should be off. If any fault LED is on, an error has been detected. Refer to "Troubleshooting" on page 359, for additional information.
6. Close and lock the disk enclosure doors.
7. If the host was shutdown to install the array, boot the host.
8. Perform an ioscan to verify that the host sees the array.

**Table 27** Normal LED Status for the Disk Enclosure

| Module | LED | Normal State |
|---|---|---|
| Front Enclosure | System Fault | Off |
| | System Power | On (green) |
| | Disk Activity | Flashing (green) when disk is being accessed. |
| | Disk Fault LED | Off |
| Power Supply | Power Supply | On (green) |
| BCC Module | Term. Power | On (green) |
| | Full Bus | Off (if split bus) On (green - if single bus) |
| | BCC Fault | Off |
| | Bus Active | On (green bus is available for use)<br>Off (Isolator chip disabled & bus not avail.) |
| | LVD | On (green bus operating in LVD mode)<br>Off (bus operating in single ended mode) |
| Fan Module | Fan | On (green) |

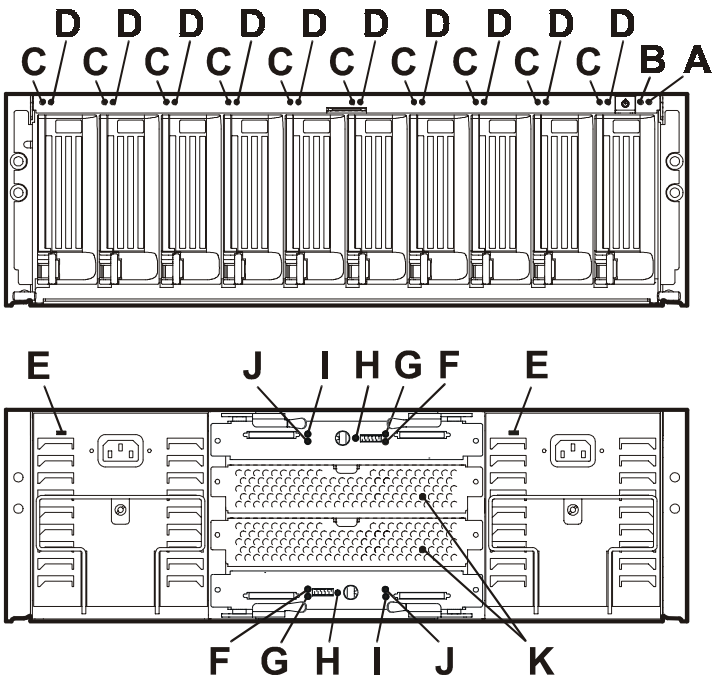

D D D D D D D D D D
C C C C C C C C C C B A
E J I H G F E
F G H I J K
A System fault LED
B System power LED
C Disk activity LED
D Disk fault LED
E Power On LED
F Term. Pwr. LED
G Full Bus LED
H BCC Fault LED
I Bus Active LED
J LVD LED
K Fan Fault LED

**Figure 77** Disk Enclosure LEDs

Table 28 Normal LED Status for Controller Enclosure

| Module | LED | Normal State |
|---|---|---|
| Controller Enclosure | Power On | On (green) |
| | Power Fault | Off |
| | Fan Fault | Off |
| | Controller Fault | Off |
| | Fast Write Cache | On (green) while data is in cache |
| Controller | Controller Power | On (green) |
| | Controller Fault | Off |
| | Heartbeat | Blink (green) |
| | Status | Green<br>There are 8 status LEDs. The number and pattern of these LEDs depend on how your system is configured. |
| Controller Battery | Fault - B | Off |
| | Full Charge - B | On (green)[1] |
| | Fault - A | Off |
| | Full Charge - A | On (green)[1] |
| Controller Power Assembly | Power 1 | On (green) |
| | Power 2 | On (green) |
| Controller Fan Assembly | Fan Power | On (green) |
| | Fan Fault | Off |

1. Both Full Charge A and Full Charge B LEDs are ON after batteries are fully charged. The LEDs flash while charging is in progress, and remain on when charging is complete.

Any time power to the disk array controller enclosure is cycled, the BBU will recharge for 15 minutes. During this time, the Full Charge LEDs will flash.

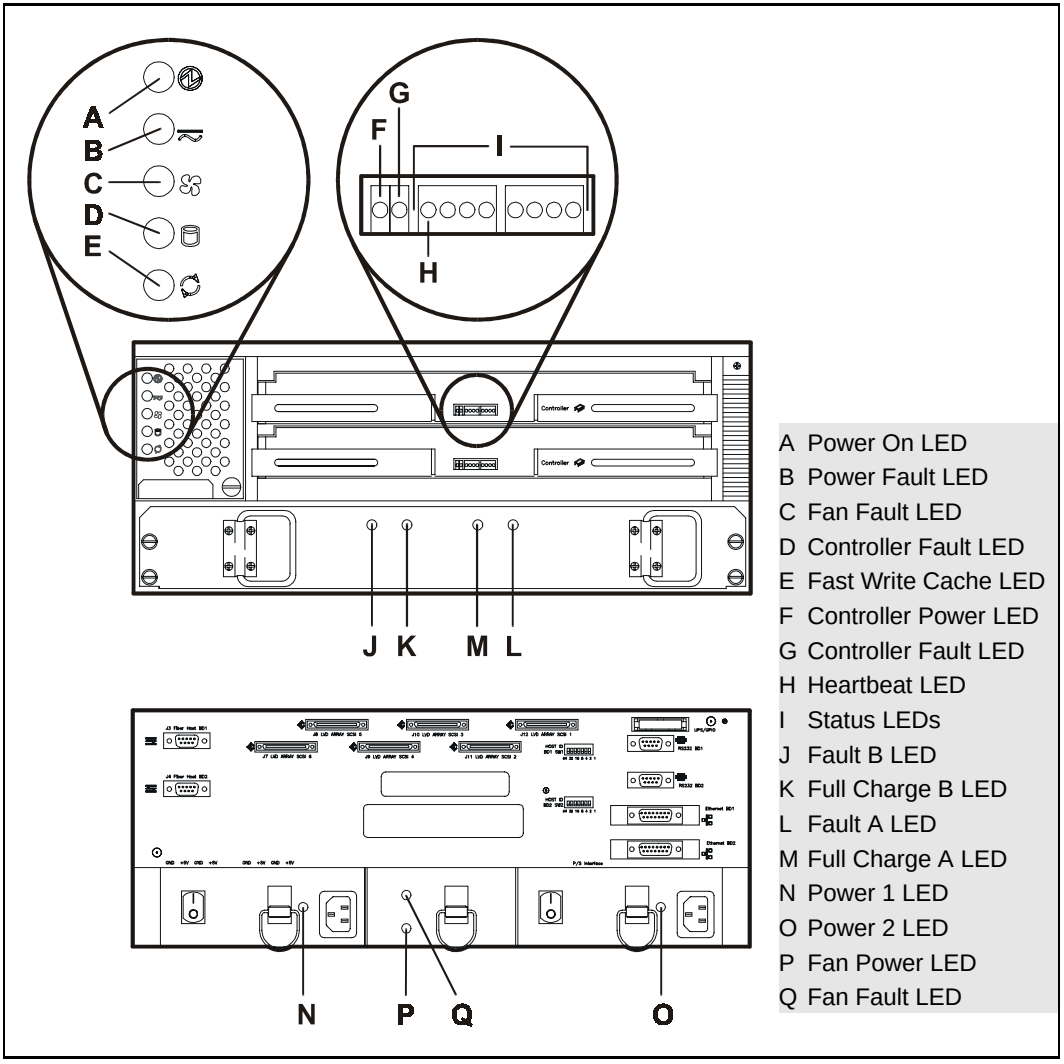



**Figure 78** Controller Enclosure LEDs

### Powering Down the Array

When powering down the disk array, the controller enclosure should be powered down before the disk enclosures.

To power down the disk array:

1. Stop all I/Os from the host to the disk array.
2. Wait for the Fast Write Cache LED to go off, indicating that all data in cache has been written to the disks.
3. Power down the controller enclosure.
4. Power down the disk enclosures.

**Note** If necessary, power to all components of the disk array can be shut off simultaneously using a single cabinet power switch or circuit breaker. In this situation, stop all I/O to the disk array and then wait two minutes for all cache data to be written to disk before shutting off power. The Fast Write Cache LED will remain off when all cache data is written.

# Verifying Disk Array Connection

## On Windows NT and Windows 2000

The *HP Storage Manager 60* software is used to verify that the disk array is visible to the Windows host. See the *HP Storage Manager 60 User's Guide* for instructions on installing and using the *HP Storage Manager 60* software.

## On HP-UX

To verify that the Disk Array FC60 is visible to the HP-UX host, run an ioscan by typing the following:

**ioscan -fn**

An output will be displayed similar to that in Figure 79. The sample ioscan shown in the figure includes entries for both disk array controllers: one at loop ID 6, and the other at loop ID 7 (loop IDs are determined by the Fibre Channel Host ID switch setting on the back of the controller enclosure). The virtual SCSI bus entries for both controllers are also shown.

**Note** The State of the each controller and its associated LUNs should be CLAIMED. Any other condition indicates that the host is having difficulty accessing the controller. If the State is not CLAIMED, refer to "Troubleshooting" on page 359 for help in identifying the problem.

```
Class     I    H/W Path           Driver   State     H/W Type   Description
=============================================================================
target    0    8/8.8.0.6.0.0      tgt      CLAIMED   DEVICE
disk      0    8/8.8.0.6.0.0.0    sdisk    CLAIMED   DEVICE     HP A5277A      1
                                  /dev/dsk/c0t0d0  /dev/rdsk/c0t0d0
disk      1    8/8.8.0.6.0.0.1    sdisk    CLAIMED   DEVICE     HP A5277A
                                  /dev/dsk/c0t0d1  /dev/rdsk/c0t0d1
disk      2    8/8.8.0.6.0.1.0    sdisk    CLAIMED   DEVICE     HP A5277A      2
                                  /dev/dsk/c0t1d0  /dev/rdsk/c0t1d0
disk      3    8/8.8.0.6.0.2.0    sdisk    CLAIMED   DEVICE     HP A5277A      3
                                  /dev/dsk/c0t2d0  /dev/rdsk/c0t2d0
disk      4    8/8.8.0.6.0.3.0    sdisk    CLAIMED   DEVICE     HP A5277A      4
                                  /dev/dsk/c0t3d   /dev/rdsk/c0t3d0
target    4    8/8.8.0.7.0.0      tgt      CLAIMED   DEVICE
disk      5    8/8.8.0.7.0.0.0    sdisk    CLAIMED   DEVICE     HP A5277A
                                  /dev/dsk/c1t0d0  /dev/rdsk/c1t0d0
disk      6    8/8.8.0.7.0.0.1    sdisk    CLAIMED   DEVICE     HP A5277A
                                  /dev/dsk/c1t0d1  /dev/rdsk/c1t0d1
disk      7    8/8.8.0.7.0.1.0    sdisk    CLAIMED   DEVICE     HP A5277A
                                  /dev/dsk/c1t1d0  /dev/rdsk/c1t1d0
disk      8    8/8.8.0.7.0.2.0    sdisk    CLAIMED   DEVICE     HP A5277A
                                  /dev/dsk/c1t2d0  /dev/rdsk/c1t2d0
disk      9    8/8.8.0.6.0.3.0    sdisk    CLAIMED   DEVICE     HP A5277A
                                  /dev/dsk/c1t3d0  /dev/rdsk/c1t3d0
target    8    8/8.8.0.255.0.6    tgt      CLAIMED   DEVICE                    5
ctl       0    8/8.8.0.255.0.6.0  sctl     CLAIMED   DEVICE     HP A5277A
                                  /dev/rscsi/c2t6d0
target    9    8/8.8.0.255.0.7    tgt      CLAIMED   DEVICE                    6
ctl       1    8/8.8.0.255.0.7.0  sctl     CLAIMED   DEVICE     HP A5277A
                                  /dev/rscsi/c2t7d0
```

**1** **Entry for virtual SCSI bus 0 on controller module with loop ID 6 (VSA)**

**2** **Entry for virtual SCSI bus 1 on controller module with loop ID 6 (VSA)**

**3** **Entry for virtual SCSI bus 2 on controller module with loop ID 6 (VSA)**

**4** **Entry for virtual SCSI bus 3 on controller module with loop ID 6 (VSA)**

**5** **Entry for disk array controller module with loop ID 6 (PDA)**

**6** **Entry for disk array controller module with loop ID 7 (PDA)**

**Figure 79** IOSCAN Output

## Interpreting Hardware Paths

Each component on the disk array is identified by a unique hardware path. The interpretation of the hardware path differs depending on the type of addressing used to access the component. Two types of addressing are used with the Disk Array FC60:

- Peripheral Device Addressing (PDA) - used to address the disk array controller modules.
- Volume Set Addressing (VSA) - used to address the disk array LUNs

Because each controller module is identified by a unique loop address, the disk array is identified by two hardware paths.

### Peripheral Device Addressing

Peripheral device addressing (PDA) is used to identify the disk array controller modules. A disk array controller module can be identified in the ioscan output as an entry with a driver type of `sctl` and a description of HP A5277A. Figure 80 illustrates the interpretation of the fields in the hardware path when using PDA.

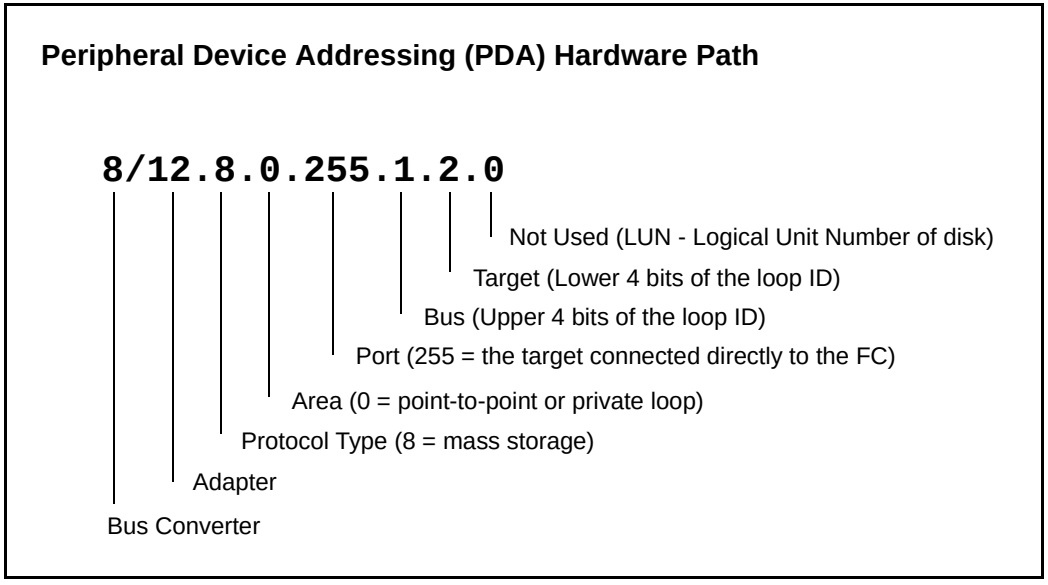

Peripheral Device Addressing (PDA) Hardware Path
8/12.8.0.255.1.2.0
Not Used (LUN - Logical Unit Number of disk)
Target (Lower 4 bits of the loop ID)
Bus (Upper 4 bits of the loop ID)
Port (255 = the target connected directly to the FC)
Area (0 = point-to-point or private loop)
Protocol Type (8 = mass storage)
Adapter
Bus Converter

**Figure 80** Peripheral Device Addressing Hardware Path

The port value will always be 255 when using PDA. The loop address, Fibre Channel Host ID of the disk array controller module (two address possible, one for controller module A and one for module B), is encoded in the Bus and Target segments of the hardware path. For example, in the hardware path shown in Figure 80, the Bus value of 1 and Target value of 2 result in a loop ID of 18:

| BUS | TARGET | LOOP ID |
|---|---|---|
| 0 0 0 1 | 0 0 1 0 | = 18 |

### Volume Set Addressing

Volume set addressing (VSA) is used to address disk array LUNs. A disk array LUN can be identified in the ioscan output as an entry with a driver type of `sdisk` and a description of HP A5277A. The LUN number is encoded in the Volume Number segments of the hardware path.

Figure 81 illustrates the interpretation of the fields in the hardware path when using VSA.

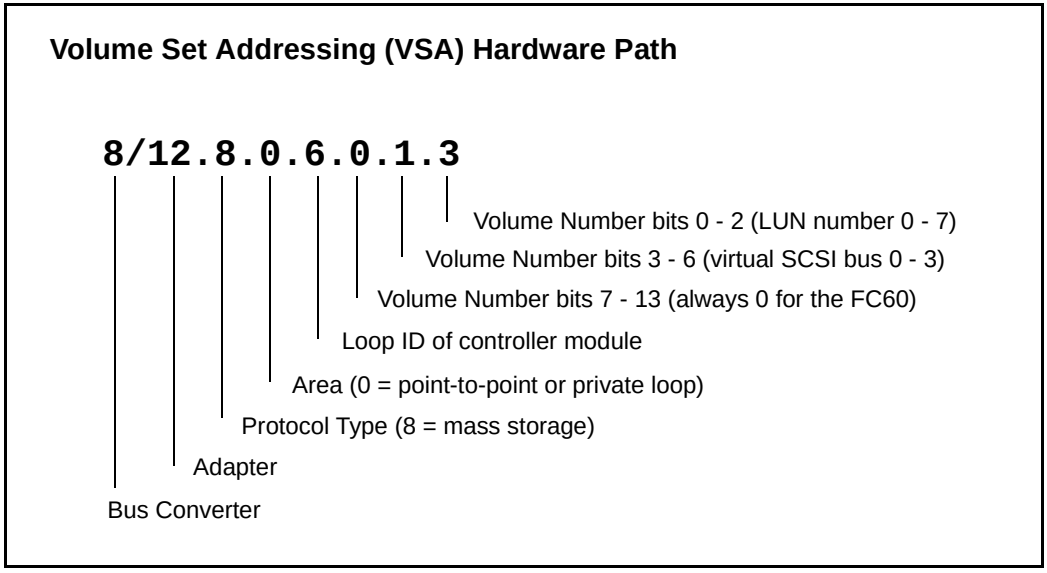

Volume Set Addressing (VSA) Hardware Path
8/12.8.0.6.0.1.3
Volume Number bits 0 - 2 (LUN number 0 - 7)
Volume Number bits 3 - 6 (virtual SCSI bus 0 - 3)
Volume Number bits 7 - 13 (always 0 for the FC60)
Loop ID of controller module
Area (0 = point-to-point or private loop)
Protocol Type (8 = mass storage)
Adapter
Bus Converter

**Figure 81** Volume Set Addressing Hardware Path

VSA is an enhancement that increases the number of LUNs that can be addressed on a fibre channel disk array to 16382 ($2^{14}$). This compares with the 8 LUN limit imposed by PDA. The HP SureStore E Disk Array FC60 requires that 32 LUNs (0 - 31) be addressable.

To implement VSA, the fibre channel driver creates four virtual SCSI busses, each capable of supporting up to eight LUNs. Each virtual SCSI bus supports the following LUNs:

| Virtual SCSI Bus | Supported LUNs |
|---|---|
| 0 | 0 - 7 |
| 1 | 8 - 15 |
| 2 | 16 - 23 |
| 3 | 24 - 31 |

The LUNs for each virtual SCSI bus are numbered 0 - 7, but reflect the LUN numbers assigned to each bus. For example, LUN 0 on virtual SCSI bus 2 is really LUN 16 on the disk array.

The virtual SCSI bus is identified by the next-to-last segment of the hardware path. For example, in the following ioscan entry, the value identifying virtual SCSI bus 1 is shown in bold. Figure 79 shows all four virtual SCSI bus entries for a disk array.

```
disk      4    8/8.8.0.0.0.1.0    sdisk    CLAIMED   DEVICE     HP A5277A
                                  /dev/dsk/c9t1d0   /dev/rdsk/c9t1d0
```

## Why Are There ioscan Entries for LUNs I Didn't Bind?

By default, an ioscan entry is made for LUN 0 of each virtual SCSI bus. This does not imply that there is an actual LUN 0 bound on the disk array. To determine if there is a LUN on the bus, use the `diskinfo` command.

For example, the sample entry above is for LUN 0 on virtual SCSI bus 1, or LUN 8. To determine if there is a LUN on virtual SCSI bus 1, perform a `diskinfo` command on the raw device file for that entry as follows:

**diskinfo /dev/rdsk/c9t1d0**

The following information is returned:

```
SCSI describe of dev/rdsk/c9t1d0
vendor: hp
product: id
type: direct access
size: 10272kbytes
bytes per sector: 512
```

If the LUN exists, the size will reflect the capacity of the LUN. If the size returned is 0 kbytes, there is no LUN 0 created for that logical SCSI bus.

### Determining LUN numbers from the hardware path

LUN numbers can be determined by using the last three segments of the VSA hardware path, which represent the 14-bit volume number. The bits of the volume number are encoded as shown in Figure 82.

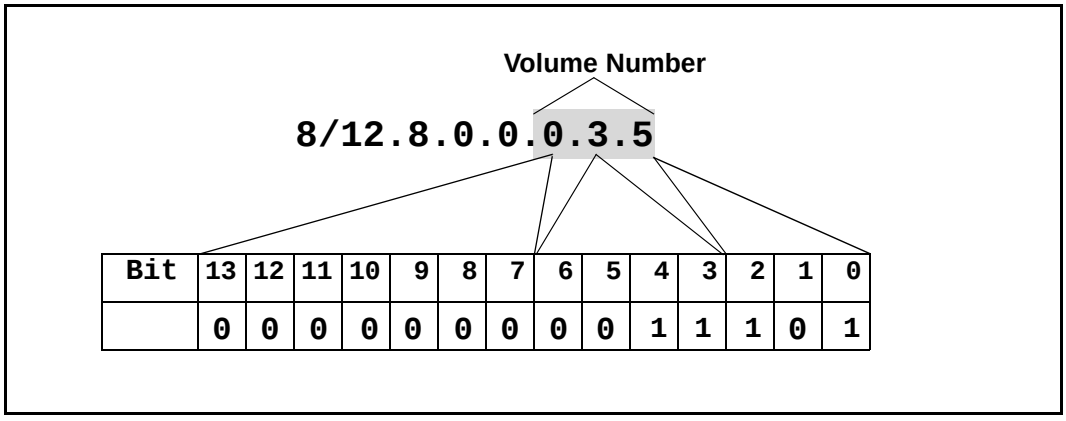



**Figure 82** VSA Volume Number Encoding

For example, if the hardware path is 8/12.8.0.0.0.3.5, the LUN number would be 29 as shown below.

```
0000000 0011 101 = 00 0000 0001 1101 = 29
```

A quick way to determine the LUN number is to multiply the value of the next-to-last segment times 8, and add the result to the last segment value. Using the above example, the LUN number is computed as follows:

```
(3 x 8) + (5) = 29
```

# Installing the Disk Array FC60 Software (HP-UX Only)

Once the disk array hardware is installed and operating, the disk array management software, diagnostic tools, and system patches must be installed. The software required for the Disk Array FC60 is distributed on the *HP-UX Support Plus CD-ROM* (9909 or later).

## System Requirements

The following host system requirements must be met to install and use the Array Manager 60 utilities successfully.

Table 29 System Requirements

| | |
|---|---|
| Hardware platform | HP 9000 Series, K-class, D-class, V-class, N-class, T-600 |
| HP-UX version | HP-UX 11.0, 11.11, and 10.20 |

## Verifying the Operating System

The Disk Array FC60 is supported on the following operating systems versions:

- HP-UX 11.0, 11.11, and 10.20

Before installing the Disk Array FC60 system software, verify that you are running the required operating system version. To identify the operating system version, type:

**uname -a**

A response similar to the following should be displayed indicating the version of HP-UX the host is running:

```
HP-UX myhost B.11.0.0 C 9000/750 1234567890 n-user license
```

**Note** Controller firmware HP08 is not supported on HP-UX 11.11. If the disk array is being installed on a host running HP-UX 11.11, it will be necessary to downgrade the controller firmware to HP03 after installing the software.

## Installing the Disk Array FC60 Software

The following HP-UX patches are required for the HP SureStore E Disk Array FC60.

- **HP-UX 10.20**: PHCO_22627 and PHSS_22846
- **HP-UX 11.0**: PHCO_22628 and PHSS_22847

The required patches can be downloaded from the following web sites:

http://us-support2.external.hp.com/index.html/ (External web site)
ftp://hpatlse.atl.hp.com/hp-ux_patches/ (Internal web site)

**Note** The patches are not currently included on the *HP-UX Support Plus CD-ROM*. They must be downloaded from the indicated web sites.

To install the software:

1. Download and install the required patches.
2. Execute the following command to verify that the required patches are installed:

**swlist**

3. Execute the following command to create the required device files (this is not required if the system was re-booted):

   **insf -e**

4. Run the Array Manager 60 amdsp command to re-scan the system for the array:

   **amdsp -R**

## Downgrading the Disk Array Firmware for HP-UX 11.11 Hosts

Controller firmware HP08 is not supported on HP-UX 11.11. If the disk array is being installed on a host running HP-UX 11.11, it will be necessary to downgrade the controller firmware to HP03 after installing the software.

1. Identify the version of firmware on the disk array controllers by typing:

   **amdload -i ArrayID**

2. If the controller firmware is HP08, contact HP support for assistance in downgrading the controller firmware.

# Configuring the Disk Array

## HP-UX

After installing the disk array software, the following steps must be performed to configure the disk array for operation. These steps should be performed by the service-trained installer with assistance from the customer where appropriate.

### Step 1. Determine the Disk Array ArrayID

The ArrayID is used to identify the disk array when performing the remaining tasks, so the first step is to determine the ArrayID.

To display the ArrayID, type:

**amdsp -i**

Record the ArrayID (its S/N) for use in performing the remaining steps.

### Step 2. Assign an Alias Name (Optional)

To simplify managing the disk array(s), an alias name can be assigned to each disk array. Once assigned, the alias can be used in place of the ArrayID when performing the remaining steps. An alias can include upper case characters, numbers, pound sign (#), period (.), and underscore (_). All other characters are invalid.

To assign an alias, type:

**ammgr -D <ArrayAlias> <ArrayID>**

- For more information, see "Assigning an Alias to the Disk Array" on page 297
- To use SAM, see "Assigning an Alias to the Disk Array" on page 264

## Step 3. Reformat Disk Array Media

***CAUTION*** *This step will destroy all data on the disk array and remove any LUN structure that has been created. If there is data on the disk array, make sure it is backed up before performing this step.*

*If a LUN structure has been created on the disk array either at the factory (A5277AZ) or by a reseller, do not perform steps 3, 4, and 5. Go to step 6 to continue configuration of the disk array.*

The disk media should be reformatted to its factory default configuration. This operation will destroy any data on the disk array.

To initialize the disk media, type:

**amutil -syswipe <ArrayID>**

**Note** As the syswipe operation executes, watch the activity LEDs on the disk modules. Each LED should illuminate briefly, starting with the first disk on the first channel, and continiung in a orderly sequence through all the disks. If the LED on any disk remains on noticeably longer than the others, it may indicate that the syswipe operation has encountered a problem. In this case, retry the syswipe operation. Do not continue wth the next step until the syswipe has completed successfully.

## Step 4. Set Both Controllers to Active State

The preceding step sets controller B to the passive states. To reset both controllers to the active state, type:

**ammgr -c AA <ArrayID>**

## Step 5. Replace LUN 0

LUN 0 was created solely to allow the host to communicate with the disk array when it is first powered on. It should be replaced with a usable LUN. If not replaced, a substantial amount of storage will be wasted.

To replace LUN 0, type:

**amcfg -R <cntrlr>:0 -d <channel:ID>,<channel:ID>..... -r <RAIDlevel> <options> <ArrayID>**

For example, the following command replaces LUN 0 with a RAID 0/1 LUN using four disks:

**amcfg -R A:0 -d 1:0,2:0,3:0,4:0 -r 1 <options> <ArrayID>**

- For a complete explanation of the command syntax and parameters, refer to the amcfg man page, or see "Replacing a LUN" on page 294.

## Step 6. Disable WCE On All Disks

For optimum data protection, WCE (Write Cache Enable) should be disabled on all disks in the array.

To disable WCE on all disks, type:

**amutil -w <ArrayID>**

**Note** To ensure optimum protection against data loss, it is recommended that Write Cache Enable be disabled on all disks in the array. Disabling disk WCE will impact disk array performance. However, it reduces the potential for data loss during a power loss.

## Step 7. Set Disk Array Time and Date

The disk array maintains its own time and date internally. These values are used to write time stamps on log entries. The disk array settings should be synchronized with the time

settings on the host to ensure valid time stamps. This ensures that any information created by the disk array such as log entries reflect the proper time they occurred.

To set the controller date and time, type:

**ammgr -t <ArrayID>**

## Step 8. Bind LUNs

If a LUN structure has been created on the disk array either at the factory (A5277AZ) or by a reseller, it may not be necessary to perform this step. If no LUN structure has been created, only LUN 0 will exist. When performing this step, consult with the customer to ensure that the desired LUN configuration is created on the disk array.

To bind a LUN, type:

**amcfg -L <cntrlr>:<LUN> -d <channel:ID>,<channel:ID>,..... -r <RAIDlevel> <options> <ArrayID>**

For example, the following command creates LUN 1 as RAID 5 using six disks:

**amcfg -L A:1 -d 1:1,2:10,3:10,4:1,5:1,6:1 -r 5 <options> <ArrayID>**

- For more information, see "Binding a LUN" on page 289
- To use SAM, see "Binding a LUN" on page 267

## Step 9. Add Global Hot Spares (Optional)

Hot spares provide an additional level of data protection. A hot spare automatically replaces a failed disk, restoring redundancy and protecting against a second disk failure. For maximum protection against disk failure, it is recommended that one hot spare be created per channel.

To add a global hot spare, type:

**ammgr -h channel:ID <ArrayID>**

For example, the following command creates a global hot spare:

**ammgr -h 4:3 <ArrayID>**

- For more information, see "Adding a Global Hot Spare" on page 296
- To use SAM, see "Adding a Global Hot Spare" on page 273

### Step 10. Install Special Device Files

After binding LUNs, you must install special device files on the LUNs. This makes the LUNs usable by the operating system.

To install the special device files, type:

**insf -e**

### Step 11. Check Disk Array Status

The final step is to display the disk array status to ensure that all features are enabled and that the array is working properly.

To display disk array status, type:

**amdsp -a <ArrayID>**

See Figure 88 on page 284 for a sample status display output that identifies the important fields and the normal values.

## Windows NT and Windows 2000

Perform the following steps to configure the disk for operation on a Windows NT or Windows 2000 host. Refer to the *HP Storage Manager 60 User's Guide* for detailed instructions on performing each of these tasks.

1. Add the disk array to the SM60 management topology.
2. Set up any alert notifications.
3. Rename the disk array.
4. Create the desired volume structure. Replace the default 10 Mbyte volume if necessary.
5. Add hot spares as required.

6. Set up storage partitions if this premium feature is enabled.
7. Set the disk array controller clocks

# Using the Disk Array FC60 as a Boot Device (HP-UX Only)

The Disk Array FC60 is supported for use as boot device on the following HP 9000 Servers running HP-UX 11.0:

- K-Class
- D-Class
- N-Class
- L-Class

**Note** Not all levels of server PDC (processor dependent code) support Fibre Channel boot. Before performing the following procedure, ensure that the level of PDC on the server supports booting from a Fibre Channel device.

To serve as a boot device, the Disk Array FC60 must have a suitable LUN created on it. If the disk array has a LUN structure on it, an available LUN can be used for the boot disk, or another LUN can be created. This applies to a disk array that has already been in use, and when installing a new factory assembled disk array (A5277AZ).

When installing a field racked disk array,(A5277A), it will be necessary to create a LUN before the disk array can be used as a boot device.

# Solving Common Installation Problems

**Problem.** When performing an ioscan, the host sees the disk array controllers but none of the disks or LUNs in the disk array.

**Solution.** This is typically caused by powering on the disk array controller enclosure before powering on the disk enclosure(s). Turn off power to all disk array components, power on the disk enclosures and wait for the disks to spin up, then power on the controller enclosure.

**Problem.** When installing a factory-assembled disk array, Array Manager 60 does not see the array. The amdsp -R command does not detect the array, although an ioscan displays the disk array LUNs.

**Solution.** This can result if any of the disk array enclosures are not powered on, or they have been disconnected or removed from the cabinet. Factory-assembled disk arrays have an entire LUN structure created on them. If any of the disk enclosures are missing when the disk array is powered on, the LUNs will appear but they will all have a size of zero bytes. Make sure all disk enclosures are present and powered on before the controller enclosure is powered on.

**Problem.** The ioscan results shows extra LUNs on the disk array. There are several LUN 0 entries for the disk array that I did not bind.

**Solution.** The volume set addressing used by the disk array creates virtual SCSI busses to support addressing up to 32 LUNs. These LUNs do not really exist on the disk array, but are displayed in the ioscan output.

For more information, see "Volume Set Addressing" on page 209.

**Problem.** The Array Manager 60 commands are not working. When I attempt to execute a command I get a message indicating that AM60Srvr is not running.

**Solution.** If the AM60Srvr daemon has stopped, no Array Manager 60 commands can execute. AM60Srvr can be restarted using the following command:

**/opt/hparray/bin/AM60Srvr**

# Adding Disk Enclosures to Increase Capacity

Scalability is an important part of the design of the HP SureStore E Disk Array FC60. The capacity of the disk array can be increased in a variety of ways to meet growing storage needs. See "Adding Capacity to the Disk Array" on page 254 for more information on scalability options.

Adding disk array enclosure(s) is the most effective way of significantly increasing the capacity of the disk array. It is also the most involved. The following procedure describes the process of adding disk enclosures to the disk array configuration.

**Note** This procedure should be performed by properly trained personnel only.

## General Rules for Adding Disk Enclosures to the Disk Array

- **Shutting Down the Disk Array** - Adding one or more new disk enclosures requires powering off the entire disk array. This will make all data on the disk array unavailable during the expansion process.
- **Backing Up Data** -All data on the disk array should be backed up to prevent data loss should problems occur during the expansion.
- **Maintaining High Availability** - Adding disk enclosures will typically require that you convert existing disk enclosures from split-bus configuration to full-bus configuration. This may result in more than one disk in a LUN being on the same channel. This will necessitate moving disks to other enclosures following the expansion. For optimum high-availability, there should never be no more than one disk from a LUN on each channel.
- **Moving Disk Modules** - When moving disks during the expansion, it is recommended that disks be moved to the enclosure and slot that correspond to the original channel:ID of the disk. For example, if the disk was originally installed in the slot corresponding to channel 4, ID 1, it should be moved to the slot that corresponds to the same values following the expansion. Although disks can be moved to any slot, maintaining the same channel:ID simplifies LUN management.

- **Consider Adding More Than One Disk Enclosure** - Because the process of adding disk enclosures involves backing up data and powering off the disk array, you should consider adding more than one enclosure to meet your future capacity needs. This will avoid having to redo the procedure each time you add another disk enclosure. And the addition of a single enclosure provides limited flexibility for configuring LUNs on the disk array.
- **Avoid Using Split-bus and Full-bus Mode in the Disk Array** - To simplify disk array and LUN management, it recommended that you avoid mixing split-bus and full-bus disk enclosures in the same disk array.

## Step 1. Plan the Expanded Configuration

Your expansion strategy will be dictated by the amount of capacity you are adding to the disk array. This includes both the number of new enclosures you are adding, and the number of new disk modules. Careful planning will help ensure that the expansion is performed successfully.

1. Identify the original disk array configuration by performing the following tasks:
   a. Create a detailed diagram of the existing Disk Array FC60 layout.
   b. Attach a label to each Fiber Channel cable attached to the disk array, identifying the host interface and disk array controller port it connects.
   c. Identify the disks comprising each LUN on the disk array. Then attach a label to each disk in the LUN, identifying its LUN number and channel:ID. For example, LUN 0, 4:0. The channel and ID assigned to each disk are a function of the enclosure and slot the disk is installed in.
   d. Identify all global hot spare disks in the disk array. Then attach a label to each global hot spare disk, identifying the channel:ID of the disk.

**Note** A global hot spare disk should not be moved. If you intend to move a hot spare disk that is not in use, the disk must be removed from the hot spare group. This will change its role to Unassigned. After the disk has been moved to its new location, it can again be assigned the role as a global hot spare.

2. Identify the expanded disk array layout by performing the following tasks:

   a. Create a detailed diagram of the expanded HP FC60 array layout. Include all Fibre Channel and SCSI cabling connections. This diagram will serve as your configuration guide as you add the new enclosures. The Capacity Expansion Map on page 235 should assist you in identifying where disk will be moved in the new configuration.

   b. Attach a second label to each disk that is part of a LUN, identifying the disk enclosure and slot the disk will occupy in the new configuration. The disk should be moved to the enclosure and slot that corresponds to the original channel:ID of the disk.

## Step 2. Backup All Disk Array Data

***CAUTION*** *It is critical that you perform this step. Protection of user data is essential if a problem occurs during the expansion process.*

1. Stop all I/O activity to the disk array.
2. Backup all data on the Disk Array FC60.
3. Label and store the backup media in a safe place.

## Step 3. Prepare the Disk Array for Shut Down

1. Determine if any LUNs in the disk array are using a global hot spare as a result of a disk failure. If any hot spares are in use, perform the following steps:

   a. Identify and replace the failed disk that caused the use of the global hot spare.

   b. Wait for the disk array to complete the process of copying data from the hot spare to the replacement disk.

   c. Verify that the LUN is now in the OPTIMAL state.

2. Verify that the status of each LUN on the disk array is OPTIMAL. All LUNs must be in the optimal state before continuing.

***CAUTION*** *Do not proceed to the next step if any LUN is not in an optimal state and you intend to move any of the disks which comprise the LUN. Contact HP Support if the LUNs cannot be made OPTIMAL before the moving disk drives.*

3. If you intend to move any global hot spares, remove them from the hot spare group as follows:
   a. Verify that the hot spare disk to be moved is not in use.
   b. Remove the disk from the hot spare group.
4. Unmount any file systems associated with the disk array. When the file systems are no longer mounted, continue with the next step.

## Step 4. Add the New Disk Enclosures

***CAUTION*** *Use proper Electrostatic Discharge (ESD) procedures during the expansion procedure. ESD discharges can easily damage the HP FC60 Array and drive components. Wearing an ESD wrist strap grounded to the component chassis is sufficient for most environments.*

1. Power down the disk array components in the following order:
   a. Switch off power to the controller enclosure. Make sure that both power supplies have been turned off before proceeding.
   b. Switch off power to each disk enclosure.
2. Remove the Fiber Channel cables connected to the controller enclosure.
3. Disconnect all SCSI cabling from the controller enclosure and the disk enclosures.
4. Install the new disk enclosures in the rack. See "Step 4: Install the Disk Enclosure" on page 162.

5. Configure the necessary disk enclosures for full-bus operation. See "Configuration Switches" on page 176.

   Set the disk enclosure DIP switches on both BCC A and BCC B to the following settings for full-bus operation:

   sw1=1 (This switch is set to 0 for split-bus mode.)
   sw2=0
   sw3=0
   sw4=0
   sw5=0

6. Install SCSI terminators on each disk enclosure. Install a SCSI Terminator on the right-most connector on both the BBC A and BCC B cards.

***CAUTION*** *An incorrectly terminated SCSI bus can cause data loss.*

7. Install the SCSI cabling between the controller enclosure and the disk enclosures. Use your expanded configuration diagram to determine how to connect the SCSI cabling. A single SCSI cable will be used to connect each full-bus disk enclosure to the controller enclosure.

**Note** To simplify disk array management, it is recommended that the SCSI cabling be connected to the disk enclosures according to where in the rack the enclosure is installed. The disk enclosure closest to the controller enclosure should be connected to SCSI channel 0. Moving down the rack, the next disk enclosure should be connected to channel 1, the next to channel 2, and so on.

8. Set the disk Enclosure (Tray) ID switches. See "Disk Enclosure (Tray) ID Switch" on page 176.

   a. Set the Enclosure ID switches on both BCC A and BCC B cards to identify the disk enclosure. The Enclosure ID switch setting must be the same for both BCC A and BCC B.

   b. The Enclosure ID switch settings are made as follows for the disk enclosures installed.

      The enclosure connected to channel 1 should be set to 0.
      The enclosure connected to channel 2 should be set to 1.
      The enclosure connected to channel 3 should be set to 2.
      The enclosure connected to channel 4 should be set to 3.
      The enclosure connected to channel 5 should be set to 4.
      The enclosure connected to channel 6 should be set to 5.

9. Move the necessary disks from the existing disk enclosures to the new enclosures. Refer to your expanded configuration drawing and the disk labels to determine the location of each disk that must be moved.

## Step 5. Completing the Expansion

***CAUTION*** *The disk array components must be powered up in the specified sequence - disk enclosures first, followed by the controller enclosure. Failure to follow the proper sequence may result in the host not recognizing LUNs on the disk array.*

1. Ensure all power cables are connected to the controller enclosure and disk enclosures.
2. Power up the disk array in the following sequence:
   a. Power up all the disk enclosures. Wait approximately two minutes for all disk enclosure activity to complete.
   b. Power up the controller enclosure.
   c. Wait for five minutes. During this period the array will discover the new configuration. It is important that the array be allowed complete the discovery process.
3. Check for any disk faults. If a disk fault occurs, remove the disk, wait for at least 15 seconds, then reinsert the disk. The disk array controller should clear the fault on the disk once it has spun up.

**Note** If the disk fault will not clear, it may be necessary to replace the disk. If a new disk does not clear the fault, contact HP Support.

4. Add any new disks to the disk enclosures. Refer to your expanded configuration drawing for information on what slots to install the disks in.

**Note** If you are adding disks that were installed in another Disk Array FC60, make sure the disks have a status of Unassigned before installing them. If the disks are marked as part of a LUN from the original array, the disk array will try to recreate the LUN.

5. Connect the Fibre Channel cable(s) from the host to the controller enclosure. Refer to the label on the cable to ensure the proper connection is maintained. Care should be

taken not to cross the cables, as this may cause problems with applications that depend on a specific path.

6. Rescan the disk array from the host using the `ioscan -fnC disk` command. The disk array and the paths to each LUN should be displayed.

This completes the process of expanding the disk array. You can now make the capacity provided by the new disks available to the host by binding LUNs.

## Capacity Expansion Example

An example of expanding an Disk Array FC60 is shown in Figure 83. In this example, three new disk enclosures are added to a disk array with three fully loaded enclosures. The disk array is configured with five 6-disk LUNs.

The original enclosures were operating in split-bus mode, and have been reconfigured to full-bus mode. The disks have been moved from their original locations to slots with the corresponding channel:ID. This strategy achieves the desired result of having each disk in the LUN on a separate channel, and also simplifies LUN management.

**Note** In the following example the terms “LUN” and “volume group” are used interchangeably.

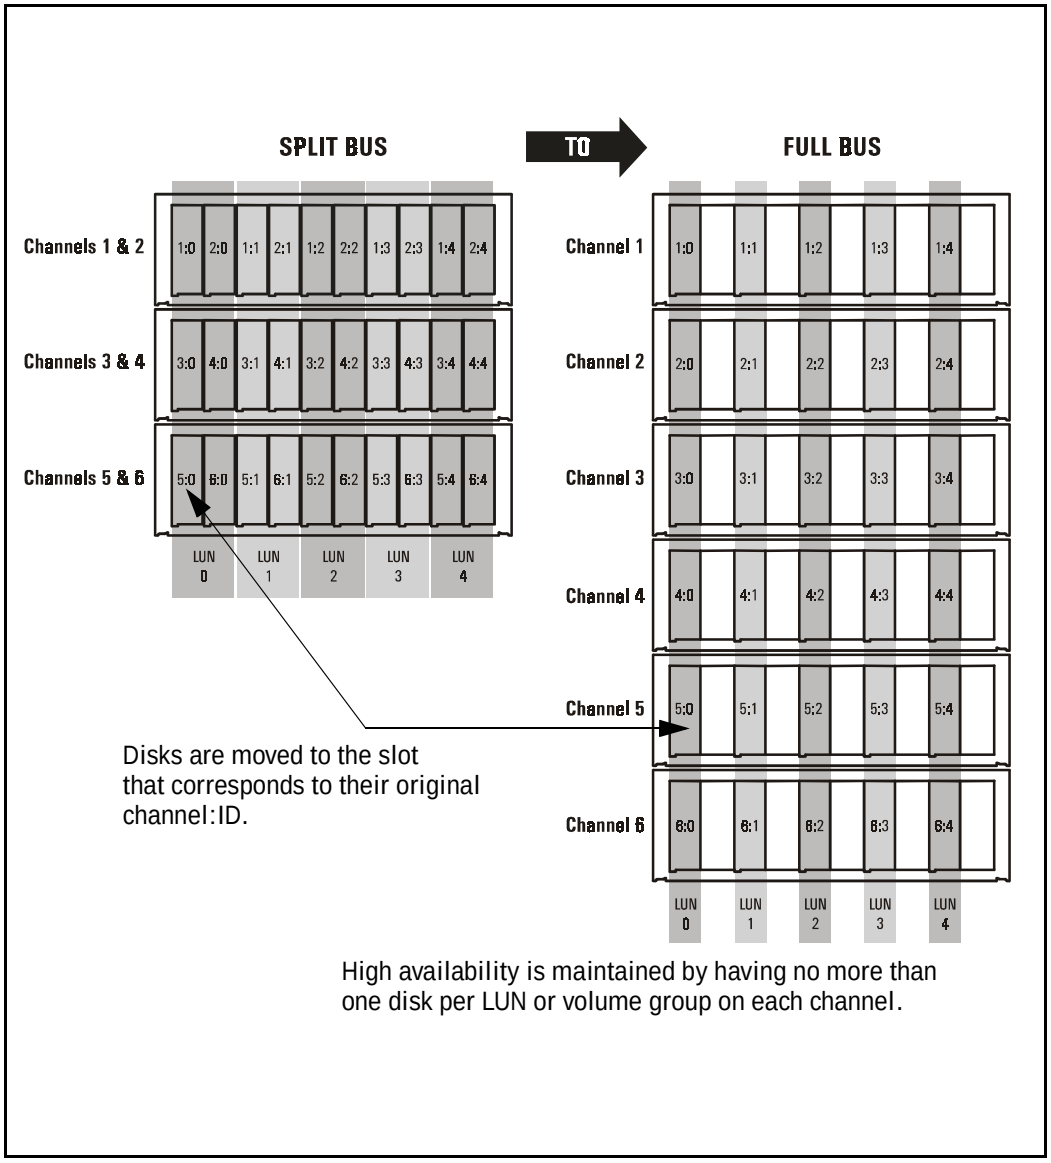

SPLIT BUS
TO
FULL BUS
Channels 1 & 2
Channels 3 & 4
Channels 5 & 6
1:0 2:0 1:1 2:1 1:2 2:2 1:3 2:3 1:4 2:4
3:0 4:0 3:1 4:1 3:2 4:2 3:3 4:3 3:4 4:4
5:0 6:0 5:1 6:1 5:2 6:2 5:3 6:3 5:4 6:4
LUN 0
LUN 1
LUN 2
LUN 3
LUN 4
Channel 1
Channel 2
Channel 3
Channel 4
Channel 5
Channel 6
1:0 1:1 1:2 1:3 1:4
2:0 2:1 2:2 2:3 2:4
3:0 3:1 3:2 3:3 3:4
4:0 4:1 4:2 4:3 4:4
5:0 5:1 5:2 5:3 5:4
6:0 6:1 6:2 6:3 6:4
LUN 0
LUN 1
LUN 2
LUN 3
LUN 4
Disks are moved to the slot that corresponds to their original channel:ID.
High availability is maintained by having no more than one disk per LUN or volume group on each channel.

**Figure 83** Capacity Expansion Example

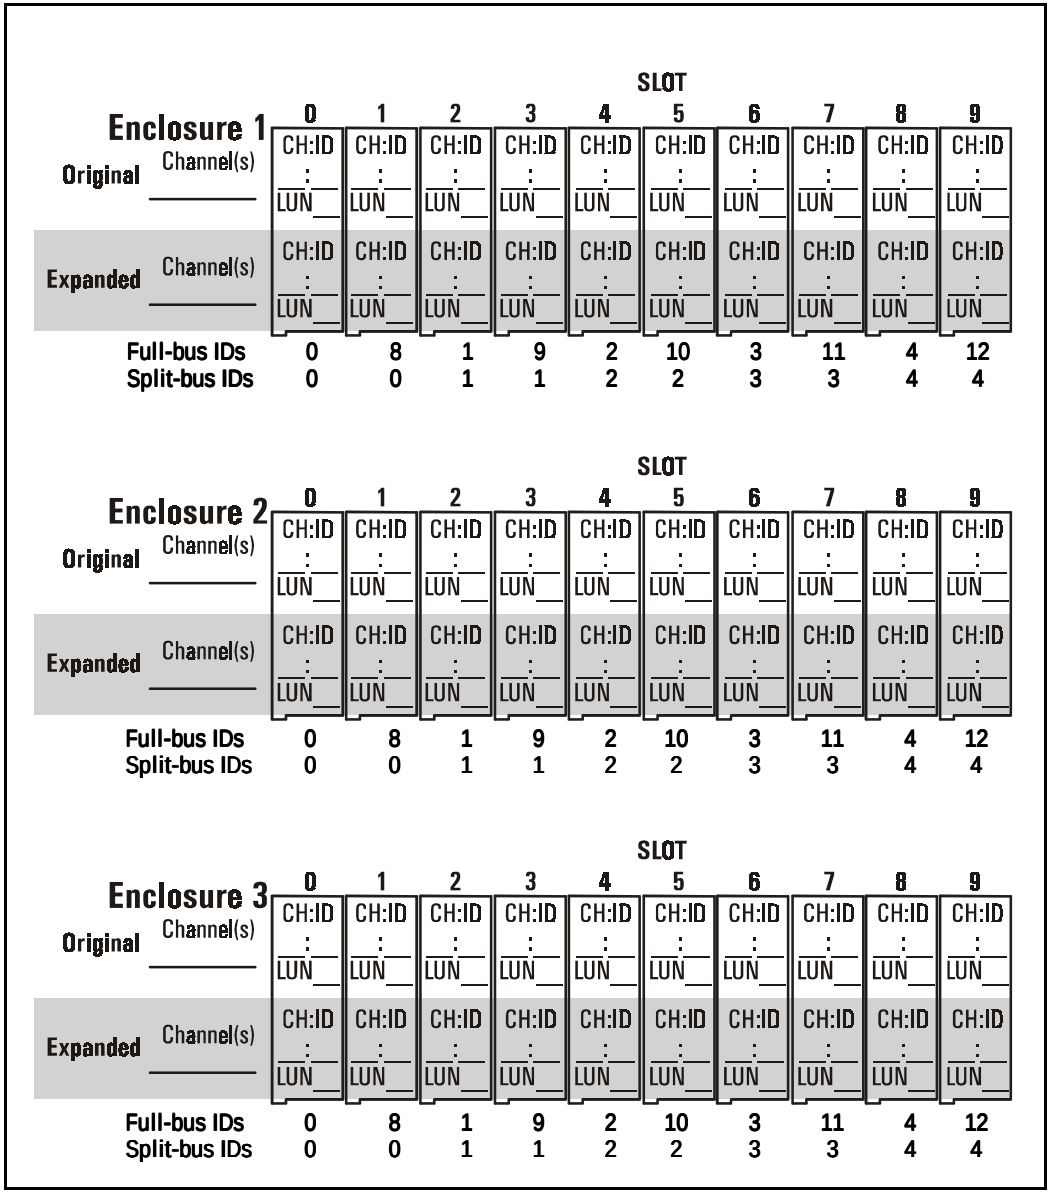
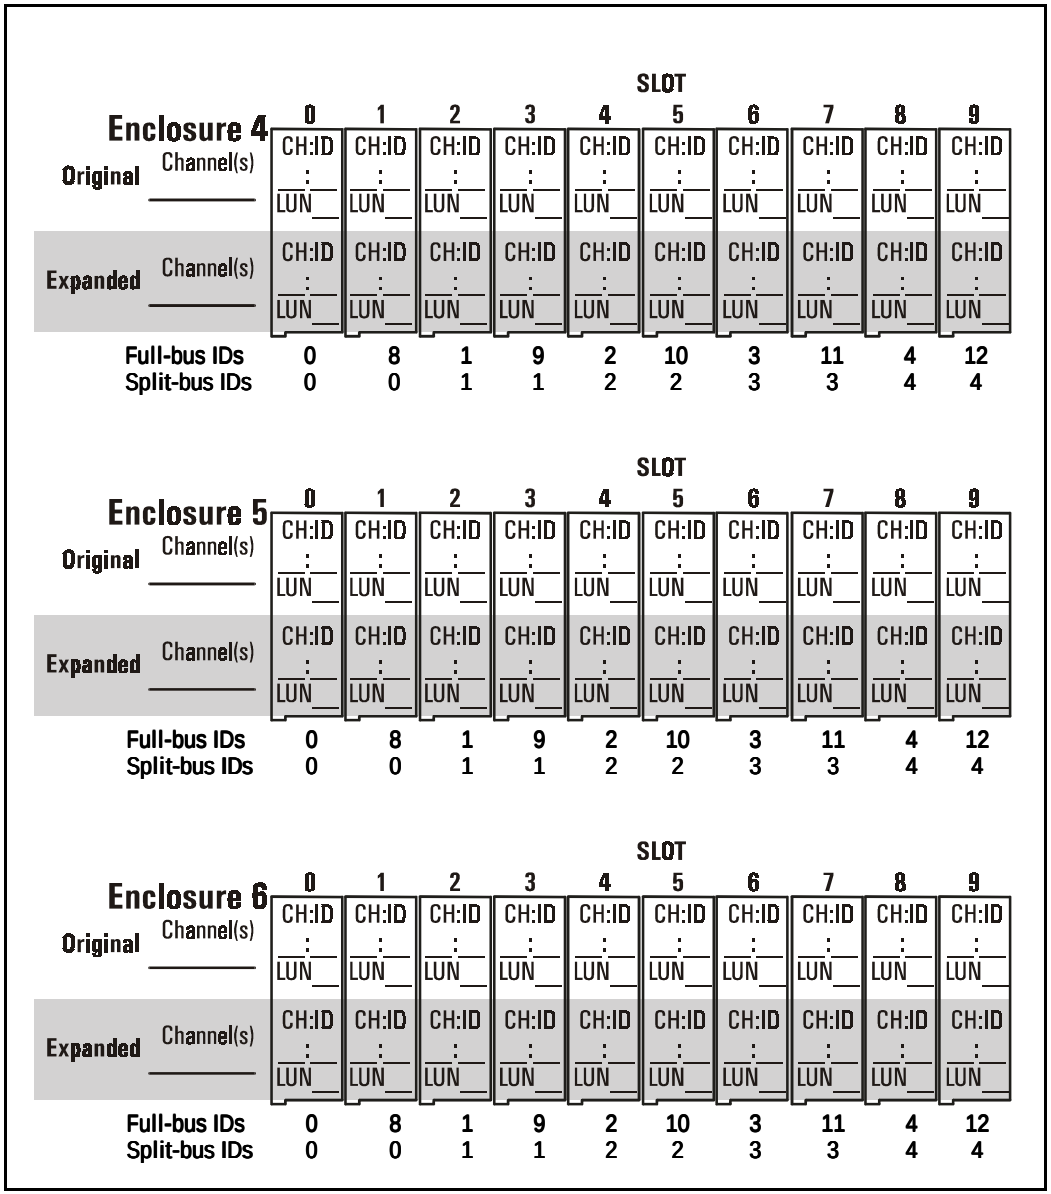

SLOT
Enclosure 1
Original
Channel(s)
Expanded
Channel(s)
0 1 2 3 4 5 6 7 8 9
CH:ID : LUN
Full-bus IDs 0 8 1 9 2 10 3 11 4 12
Split-bus IDs 0 0 1 1 2 2 3 3 4 4
SLOT
Enclosure 2
Original
Channel(s)
Expanded
Channel(s)
0 1 2 3 4 5 6 7 8 9
CH:ID : LUN
Full-bus IDs 0 8 1 9 2 10 3 11 4 12
Split-bus IDs 0 0 1 1 2 2 3 3 4 4
SLOT
Enclosure 3
Original
Channel(s)
Expanded
Channel(s)
0 1 2 3 4 5 6 7 8 9
CH:ID : LUN
Full-bus IDs 0 8 1 9 2 10 3 11 4 12
Split-bus IDs 0 0 1 1 2 2 3 3 4 4

**Figure 84** Capacity Expansion Map


SLOT
Enclosure 4
0 1 2 3 4 5 6 7 8 9
Original Channel(s)
Expanded Channel(s)
CH:ID : LUN
Full-bus IDs 0 8 1 9 2 10 3 11 4 12
Split-bus IDs 0 0 1 1 2 2 3 3 4 4
SLOT
Enclosure 5
0 1 2 3 4 5 6 7 8 9
Original Channel(s)
Expanded Channel(s)
CH:ID : LUN
Full-bus IDs 0 8 1 9 2 10 3 11 4 12
Split-bus IDs 0 0 1 1 2 2 3 3 4 4
SLOT
Enclosure 6
0 1 2 3 4 5 6 7 8 9
Original Channel(s)
Expanded Channel(s)
CH:ID : LUN
Full-bus IDs 0 8 1 9 2 10 3 11 4 12
Split-bus IDs 0 0 1 1 2 2 3 3 4 4

# 4 MANAGING THE DISK ARRAY ON HP-UX


Tools for Managing the Disk Array FC60 . . . . . 238

Installing the Array Manager 60 Software . . . . . 240

Managing Disk Array Capacity . . . . . 242

Adding Capacity to the Disk Array . . . . . 254

Upgrading Controller Cache to 512 Mbytes . . . . . 258

Managing the Disk Array Using SAM . . . . . 260

Managing the Disk Array Using Array Manager 60 . . . . . 276

Managing the Disk Array Using STM . . . . . 314

Status Conditions and Sense Code Information . . . . . 317

# Tools for Managing the Disk Array FC60

**Note** On Windows NT and Windows 2000, the disk array is managed using the *HP Storage Manager 60* software. See the *HP Storage Manager 60 User's Guide* for information on managing the disk array on Windows NT and Windows 2000.

There are three tools available for managing the Disk Array FC60 on HP-UX: the HP-UX System Administrator Manager (SAM), Array Manager 60, and Support Tools Manager (STM). The various management tasks and the tools that can be used to perform them are shown in Table 30.

## System Administration Manager (SAM)

The common tasks involved in managing the disk array can be performed using the HP-UX System Administration Manager, or SAM. This allows you to manage the disk using the same tool you use to manage your host system.

Not all tasks can be performed using SAM. In general, the more common tasks are available from SAM while the more advanced tasks are performed using Array Manager 60.

## Array Manager 60

Array Manager 60 comprises a set of utilities that provide complete capability for managing the disk array. In addition to the common tasks available through SAM, Array Manager 60 offers more advanced management capability. Although you will likely use SAM to do much of the disk array management, you will need to use Array Manager 60 to perform any tasks that are not supported by SAM.

## STM

STM, the Support Tools Manager, provides a third alternative for managing your disk array. STM can be used to perform many common disk array management tasks. See "HP-UX Diagnostic Tools" on page 345 for complete information on using STM.

**Table 30** Management Tools and Tasks

| Task | Tool | | |
|---|---|---|---|
| | **SAM** | **Array Manager 60** | **STM** |
| Checking disk array status | Yes | Yes (amdsp) | Yes |
| Managing LUNs | Yes | Yes (amcfg) | Yes |
| Managing global hot spares | Yes | Yes (ammgr) | Yes |
| Assigning an alias to the disk array | Yes | Yes (ammgr) | No |
| Managing cache memory | No | Yes (ammgr) | No |
| Managing the rebuild process | No | Yes (amutil) | No |
| Synchronizing controller date and time | No | Yes (ammgr) | No |
| Locating disk array components | Yes | Yes (amutil) | Yes |
| Performing a parity scan | No | Yes (ammgr) | Yes |
| Displaying disk array serial numbers | Yes | Yes (amdsp) | No |
| Disabling disk Write Cache Enable (WCE) | No | Yes (amutil) | No |
| Managing disk array logs | No | Yes (amlog) | Yes |

# Installing the Array Manager 60 Software

The Array Manager 60 software must be installed on the host to which the disk array is connected. The software should have been installed with disk array hardware. However, if you decide to move the disk array to a different host, you will need to install the software on the new host.

See "Installing the Disk Array FC60 Software (HP-UX Only)" on page 213 for more information.

**Note**

**Must Array Manager 60 be installed to use SAM?**
Yes. The Array Manager 60 software is required to manage the HP SureStore E Disk Array FC60 even when using SAM. When using SAM, the Array Manager 60 utilities are invoked to perform the management tasks.

## AM60Srvr Daemon

The AM60Srvr daemon is the server portion of the Array Manager 60 software. It monitors the operation and performance of all disk arrays, and services requests from clients used to manage the disk arrays. Tasks initiated from SAM or Array Manager 60 are serviced by the AM60Srvr daemon.

The AM60Srvr daemon monitors disk array performance and status, maintains disk array logs, initiates diagnostics, and allows clients to examine and change the disk array configuration.

AM60Srvr must be running to allow management of the disk array using either SAM or Array Manager 60. Because of its importance in managing the disk arrays, AM60Srvr is launched automatically when the system is booted. Host I/Os to the disk array are not dependent on AM60Srvr and are serviced regardless of whether AM60Srvr is running or not.

## Running Array Manager 60

Array Manager 60 runs as a daemon process (AM60Srvr) and is started automatically when the system is started. If for some reason Array Manager 60 stop running, an error message will be displayed when you attempt to execute any of the commands.

If you suspect that Array Manager 60 is not running, type the following command:

```
ps -ef | grep AM60Srvr
```

AM60Srvr is the Array Manager daemon and it should be displayed as running. If AM60Srvr is not running, start it using the following command:

```
/opt/hparray/bin/AM60Srvr
```

# Managing Disk Array Capacity

During installation, a LUN structure is created on the disk array. This structure may meet your initial storage requirements, but at some point additional capacity may be required. This involves adding disks and binding new LUNs.

Careful LUN planning will ensure that you achieve the desired levels of data protection and performance from your disk array. Altering the characteristics of an existing LUN such as its size and RAID level requires unbinding and rebinding the LUN — a time consuming process. With careful planning you can avoid having to reconfigure the disk array when adding disk array capacity.

## Configuring LUNs

The primary task in managing disk array capacity is creating the LUN structure you need. Factors such as capacity, RAID level, and stripe segment size must be determined for each LUN. The LUN configuration you select should be based on system environment and application needs.

The process of binding LUNs is much like adding new disks to your system. Each LUN appears to the host as a single hard disk drive. However, unlike conventional hard disks, the disk array gives you the flexibility to control the size and operating characteristics of each LUN .

When binding LUNs, consider the following:

- Any size limitations imposed by HP-UX or LVM.
- The maximum numbers of LUNs that can be created on the disk array is 32. Because LUN numbering begins at 0, the highest LUN number is 31.
- Your backup strategy. If you do unattended backup to a device such as tape, you may want to avoid creating a LUN that is larger than the capacity of the media. This allows you to backup an entire LUN without changing media.

## Selecting Disks for a LUN

When binding a LUN , you must select the disks that will be used. The capacity of the LUN is determined by the number and capacity of the disks you select, and the RAID level.

When selecting disks for a LUN , consider the following:

- To maximize high availability, select disks in different disk enclosures or on different channels. Multiple disks in the same enclosure make a RAID 5 LUN vulnerable to an enclosure failure. A RAID 1 or 0/1 LUN can survive an enclosure failure, as long as both disks of a mirrored pair are not in the same enclosure. If you attempt to select disks in the same enclosure or on the same channel when binding a LUN , you will be warned that doing so may compromise high availability.
- Select disks of the same capacity.Binding a LUN using different size disks will result in unused capacity on the larger disks. For example, binding a 4-disk LUN that includes two 9-Gbyte disks and two 18-Gbyte disks will result in a LUN with the capacity of four 9-Gbyte disks. Only 9 Gbytes of the larger disks will be used when creating the LUN .
- Consider using more disks for RAID 5LUNs . This increases both performance and storage efficiency. Because RAID 5 uses only one disk's worth of capacity for parity, more disks in the LUN will result in a greater percentage of the disk capacity used for data. For example, in a 4-disk RAID 5LUN , 75% of the capacity is used for data (3 of 4). Adding another disk to the LUN will increase the percentage of disk capacity used for data to 80% (4 of 5).

### Selecting Disks for a RAID 0/1 LUN

The order in which you select disks is important when creating a RAID 0/1 LUN. The first half of the disks you select will be the primary disks, and the second half of the disks will be the disk mirrors. To maintain data availability, the disk mirrors must be in a different enclosure than the primary disks. This applies regardless of which tool you are using to bind a RAID 0/1 LUN.

For example, in Figure 85 a 4-disk RAID 0/1 LUN is being bound using one disk enclosure on channel 1, and a second disk enclosure on channel 2. The correct order for selecting disks is 1:2, 1:3, 2:2, 2:3. This selection order creates mirrored pairs of 1:2/2:2 and 1:3/2:3. This maintains high availability because the primary disks are on channel 1, and the mirror disks are on channel 2.

Selecting disks in the incorrect order of 1:2, 2:2, 1:3, and 2:3 results in mirrored pairs of 1:2/1:3 and 2:2/2:3. This puts the primary disk and mirror disk of each pair in the same enclosure, making the LUN vulnerable to an enclosure failure.

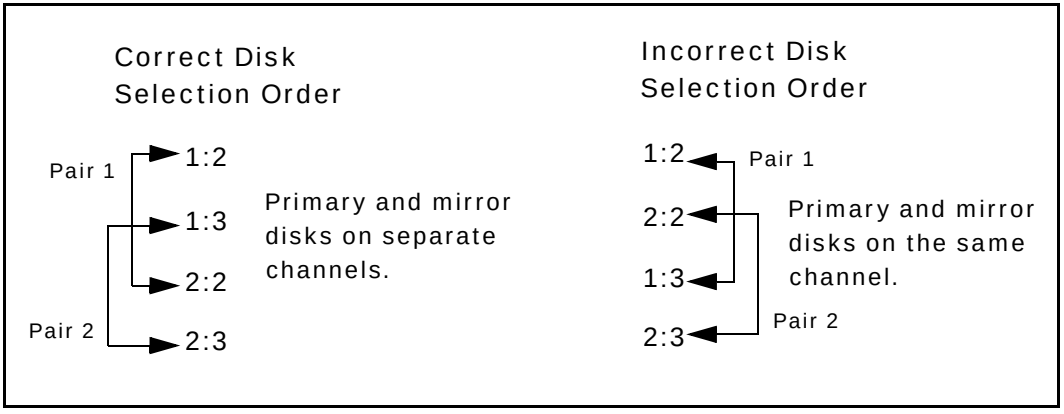



**Figure 85** Disk Selection for a RAID 0/1 LUN

### How are disk modules in the array identified?

When performing tasks such as binding LUNs , disk modules are identified by a numbered pair of the form “channel number:SCSI ID”. For example, 1:3 identifies the disk module on channel 1 with a SCSI ID of 3. Status information may also include the enclosure number and slot number.

- **Channel number** indicates the SCSI channel that the disk enclosure containing the disk module is connected to. There are six channels available (1 - 6). If the disk enclosure is using full-bus configuration, all disk modules within the enclosure will be on the same SCSI channel. If the disk enclosure is configured for split-bus operation, disks in the even-numbered slots will be on one channel, and disks in the odd-numbered slots will be on a different channel.
- **SCSI ID** identifies the ID assigned to the disk module. SCSI ID’s are determined by the slot in which the disk module is installed in the disk enclosure. If the disk enclosure is using full-bus configuration, the even-numbered slots are assigned SCSI IDs 0 - 4, and the odd-numbered slots are assigned IDs 8 - 12. (The gap in addresses is necessary for

internal management of enclosure components.) If the disk enclosure is configured for split-bus operation, the both the even-numbered slots and the odd-numbered slots are assigned IDs of 0 - 4.

- When viewing status information for the disk array, you may also see the disk **enclosure number** and **slot number** displayed. These parameters identify the physical location of the disk module within the disk array.

  The disk enclosure number is assigned during installation using the disk enclosure ID switches located on the rear of the enclosure. Each disk enclosure in the disk array must have a unique enclosure ID number. Disk enclosures should be numbered from the top down, with the enclosure mounted highest in the rack being 0.

  The slot number indicates the physical enclosure slot in which the disk module is installed. This number is related to, but is not the same as the disk SCSI ID. Slots are assigned SCSI IDs differently depending on whether the disk enclosure is configured for full or split bus operation as described above.

Figure 86 illustrates the four components involved in disk module addressing.

### Increasing LUN Capacity

The capacity of an existing LUN cannot be increased dynamically. If you need a larger LUN, you must unbind the LUN and then rebind it using more or larger disks. This requires you to back up any important data on the LUN, and then restore the data once the new, larger LUN has been bound.

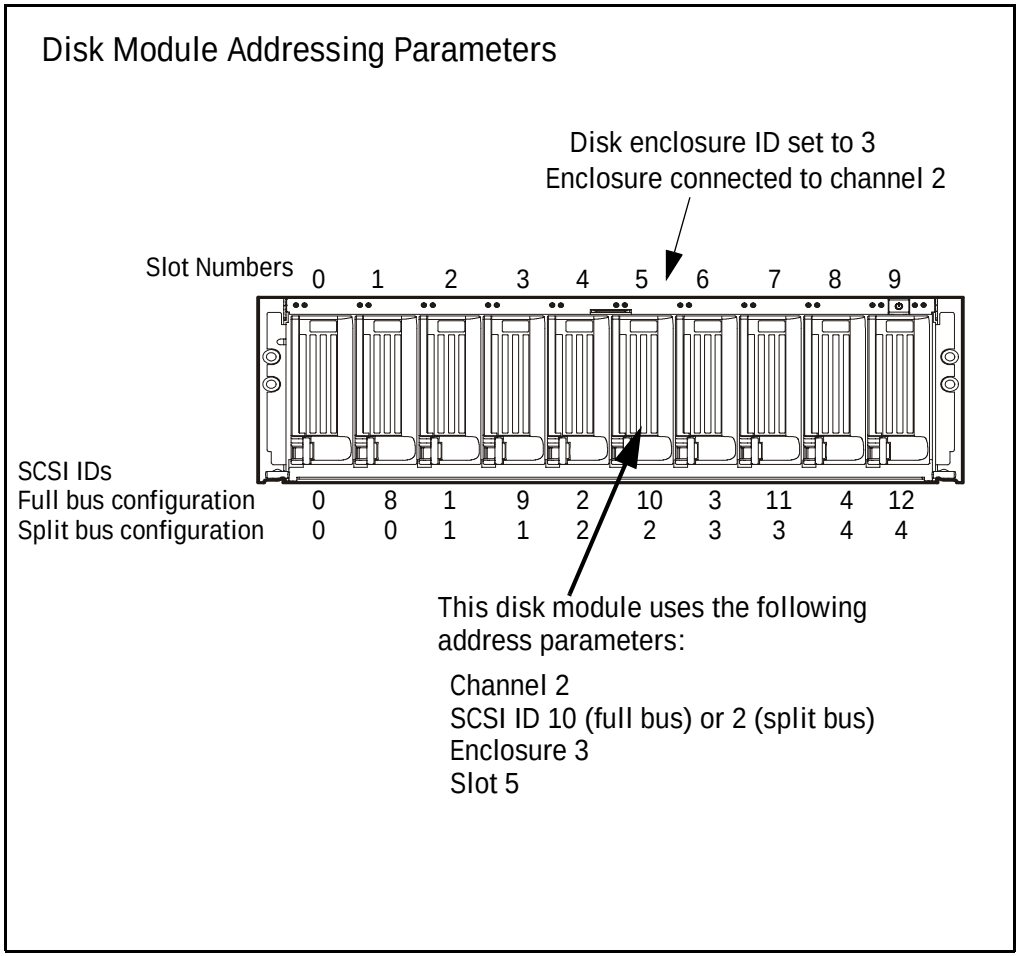

Disk Module Addressing Parameters
Disk enclosure ID set to 3
Enclosure connected to channel 2
Slot Numbers
0 1 2 3 4 5 6 7 8 9
SCSI IDs
Full bus configuration 0 8 1 9 2 10 3 11 4 12
Split bus configuration 0 0 1 1 2 2 3 3 4 4
This disk module uses the following address parameters:
Channel 2
SCSI ID 10 (full bus) or 2 (split bus)
Enclosure 3
Slot 5

**Figure 86** Disk Module Addressing Parameters

## Assigning LUN Ownership

When a LUN is bound, you must identify which disk array controller (A or B) owns the LUN. The controller that is assigned ownership serves as the primary I/O path to the LUN. The other controller serves as the secondary or alternate path to the LUN. If there is a failure in the primary I/O path and alternate links are configured, ownership of the LUN automatically switches to the alternate path, maintaining access to all data on the LUN.

When assigning LUN ownership, consider the following:

- To ensure optimum performance, LUN ownership should be balanced between controllers. This ensures that one controller is not overloaded with I/O requests, while the other is idle.
- Controller ownership can be changed on an existing LUN without impacting its operation. If you find that there is an imbalance between I/O path loads, you can change the ownership of one or more LUNs to solve the problem.

**Note** **Does the primary path selected using LVM impact LUN ownership?** Yes. The primary I/O path established using LVM defines the owning controller for the LUN. This may override the controller ownership defined when the LUN was bound. For example, if controller A was identified as the owning controller when the LUN was bound, and LVM subsequently established the primary path to the LUN through controller B, controller B becomes the owning controller.

## Selecting a RAID Level

***CAUTION*** *RAID 0 does not provide data redundancy. It should only be used in situations where high performance is more important than data protection. The failure of any disk within a RAID 0 LUN will cause the loss of all data on the LUN.. RAID 0 should only be used for non-critical data that could be lost in the event of a hardware failure.*

The RAID level you select for each LUN determines the method used to store data on the disks. Once selected, the RAID level cannot be changed dynamically. If you want to change

the RAID level used by a LUN, you must unbind the LUN and rebind it using the new RAID level.

With the exception of RAID 0, all RAID levels supported by the disk array provide protection against disk failure. However, there are differences in performance and storage efficiency between RAID levels. For more information on RAID levels and their comparative operating characteristics, see "RAID Technology" on page 47.

The two factors that will influence the RAID level you select are performance and storage efficiency.

- **Performance** - The performance characteristics differ for each RAID level. The types of applications you are running and the I/O activity associated with them should influence the RAID level you select for a LUN. The performance characteristics for each RAID level are summarized in "RAID Level Comparisons" on page 57.
- **Storage Efficiency** - the storage efficiency can range from 50% for RAID 1 and 0/1 up to > 80% for RAID 5. The higher the efficiency, the less cost per megabyte for storing your data.

## Global Hot Spares

Global hot spares provide additional protection against disk failures. The number of global hot spares you use will reflect how much protection you need. Each global hot spare you add will provide protection against the failure of a single disk. In addition to restoring hardware redundancy, a global hot spare also restores disk array performance that may be diminished while the disk array is operating in degraded mode with a failed disk.

When adding global hot spares, consider the following:

- It is recommended that you add one global hot spare per disk channel. This provides maximum protection against disk failure. A global hot spare can be used to replace any failed disk within the array, regardless of what channel it is on.
- When a failed disk is replaced, the data will be copied from the global hot spare disk to the replacement disk and the global hot spare will again be available as protection against another disk failure.

- If you choose to limit the number of global hot spares, make sure you are able to respond quickly to replace a failed disk. If an operator is always available to replace a disk, you may not need the added protection offered by multiple global hot spares.

## Setting Stripe Segment Size

Another factor you may have to consider is the stripe segment size you use for a LUN. The stripe segment size determines how much data is written to a disk before moving to the next disk in the LUN to continue writing. For example, if the stripe segment size is set to 4 Kbytes, the disk array will write 4 Kbytes of data to disk 1, then 4 Kbytes of data to disk 2, then 4 Kbytes of data to disk 3, and so on. For more information, see "Data Striping" on page 49.

When creating a LUN, the default value used for the stripe segment size is the value currently set for the cache page size (4 Kbytes or 16 Kbytes). The default cache page size is 4 Kbytes.

When setting stripe segment size, consider the following:

- Stripe segment size can affect disk array performance. The smaller the stripe segment size, the more efficient the distribution of data read or written across the stripes in the LUN. However, if the stripe segment is too small for a single I/O operation, the operation requires access to two disk. Called a stripe crossing, this action reduces performance. The optimum stripe segment size is the smallest size that will rarely force I/Os to a second disk.
- The stripe segment size can be set to any multiple of the cache page size setting. For example, if the cache page size is set to 4 Kbytes, the stripe segment size can be set to 4 Kbytes, 8 Kbytes, 16 Kbytes, 32 Kbytes, etc.

## Evaluating Performance Impact

Several disk array configuration settings have a direct impact on I/O performance of the array. When selecting a setting, you should understand how it may affect performance. Table 31 identifies the settings that impact disk array performance and what the impact is.

**Note** The LUN binding process impacts disk array performance. While a LUN is being bound, benchmarking tools should not be used to evaluate performance. Wait until all LUNS are bound to get an accurate indication of disk array performance.

Table 31 Performance Impact of Configuration Settings

| |
|---|
| **Setting:** RAID level<br>**Function:** Sets the RAID level used by the LUN.<br>**Performance Impact:** The RAID level selected impacts the entire performance profile for the LUN. Read I/O and write I/O performance are directly influenced by the RAID level. See "RAID Level Comparisons" on page 57 for more information on the performance characteristics of each RAID level. |
| **Setting:** Stripe segment size<br>**Function:** Sets the number of blocks of data the controller will write to or read from a single disk before switching to the next disk in the LUN.<br>**Performance Impact:** Optimum performance is typically achieved when the segment size matches the I/O size. In this case, only one disk is required to service an I/O, leaving the remaining disk in the LUN available for other I/Os.<br>A large segment size provides good read performance in most RAID levels. The controller will have to access fewer disks to retrieve the data, leaving the remaining drives available for other I/O operations. Large segment sizes are typically useful for applications that require high I/O throughput.<br>A small logical unit segment size is useful for most RAID 5 write applications, because the controller firmware is capable of performing group writes (writing of data simultaneously to multiple disks, while calculating the parity for the stripe, as opposed to the single-threaded read-modify-write). Small segment sizes are typically useful for applications that require large numbers of small I/Os to be processed quickly. |

Table 31 Performance Impact of Configuration Settings (cont'd)

**Setting:** Cache flush threshold (default 80%)

**Function:** Sets the level at which the disk array controller begins flushing write cache content to disk media. The setting is specified as the percentage of total available cache that can contain write data before flushing begins. The cache flush threshold can be set independently for each controller. Note that available cache is reduced by half with cache mirroring enabled. For example, if the controller has 256 Mbytes of cache and cache mirroring is enabled, only 128 Mbytes of cache is available for each controller. The remaining 128 Mbytes is being used to mirror the other controller. See Figure 87.

**Performance Impact:** This setting impacts write I/O performance by increasing or decreasing the number of write cache hits. (A "hit" is an I/O that can be serviced from cache rather than accessing the disk media.)

A higher cache threshold value increases the amount of write data that will be stored in cache before flushing to the disk is initiated. This increases the possibility that subsequent writes may be serviced from cache rather than accessing the disk, reducing the number of writes to the disk media. This improves performance for write I/Os. However, a higher value reduces the amount of cache available for reads, which may reduce the performance for read I/Os.

Lowering the cache threshold value has the opposite effect. Because less memory is used for write cache before flushing begins, the number of hits on write cache goes down with a subsequent decrease in write performance. However, read performance may improve because more memory is available for read cache.

Table 31 Performance Impact of Configuration Settings (cont'd)

| |
|---|
| **Setting:** Cache flush limit (default 100%)<br><br>**Function:** Determines how much data will remain in write cache when flushing stops. It is expressed as a percentage of the cache flush threshold. For optimum performance this value is set to 100% by default. This ensures that the entire amount of cache specified by the cache flush threshold will contain write cache data, increasing the number of write cache hits. The cache flush threshold can be set independently for each controller.<br><br>A value of 100% percent can be used for this setting because the cache flushing algorithm typically overshoots a bit in both directions. Assuming a cache flush threshold of 80% (default), the write cache contents will actually rise slightly past the threshold before flushing begins - to 82% for example. The cache flush limit of 100% (default) indicates that flushing should stop when the cache contents reaches 80%, but the flushing continues until the cache reaches a slightly lower level - 78% for example. This has the effect of maintaining a constant 80% level of data in write cache. See Figure 87.<br><br>**Performance Impact:** A high value for the flush limit increases the amount of data remaining in write cache at the end of the flush. This improves the possibility of hits on write cache. Because less data is written to the disk during each flush cycle, data flushing will occur more often, increasing disk activity. In most situations this value should be left at the default value of 100%.<br><br>A lower value reduces the amount of write data remaining in cache, thereby reducing the number of write hits. More data is flushed during each cycle, reducing the number of flushes. There will be fewer cache flush operations, but each one will last longer. |
| **Setting:** Cache page size<br><br>**Function:** Sets the number of blocks transferred into cache by the controller at one time.<br><br>**Performance Impact:** A larger cache block size increases the probability of data near the accessed block also being available in cache. However, a larger cache block size also fills up the cache sooner. A small cache block size should be used for systems that require transaction-processing requests, or I/O streams that are typically small and random. Larger cache blocks should be used for large I/O, sequential, high bandwidth applications. |

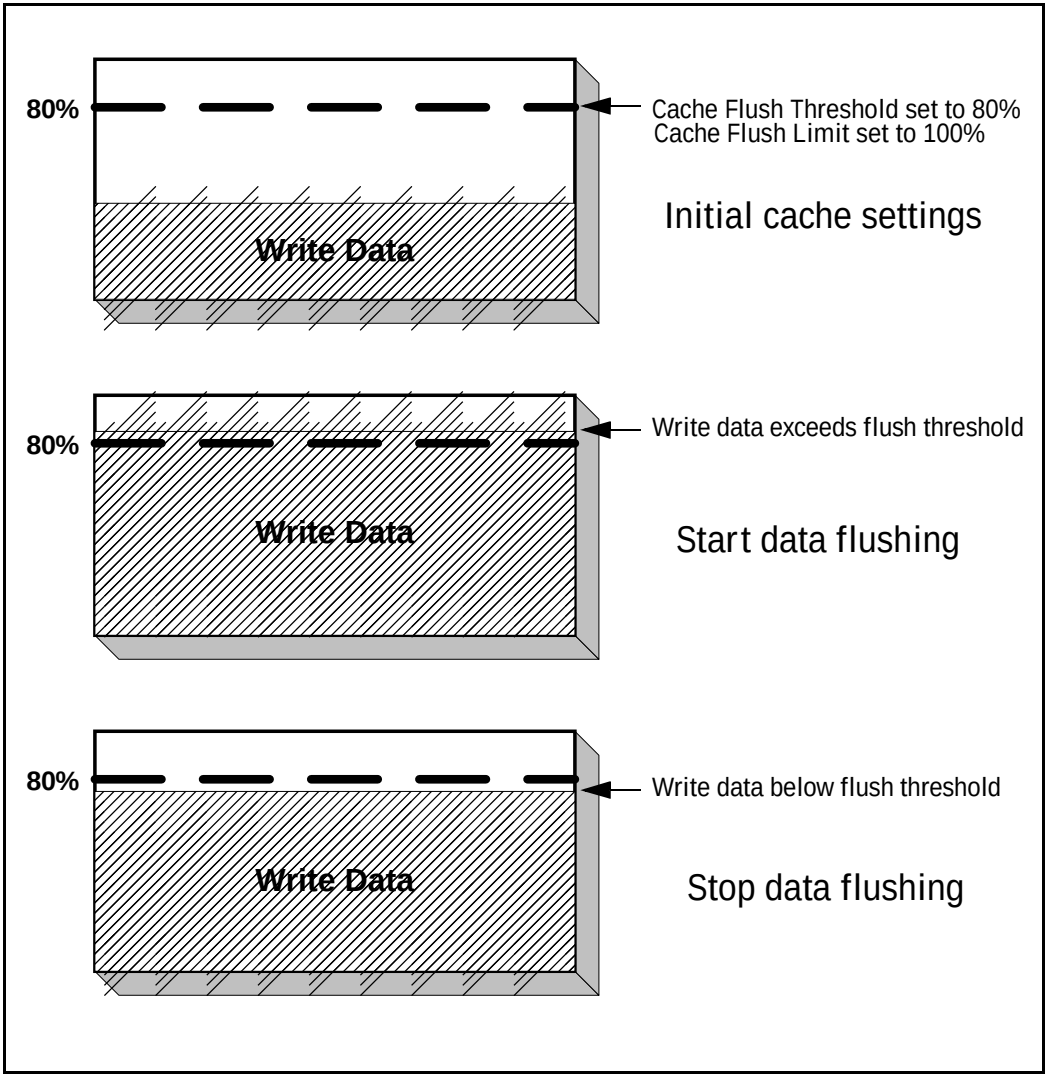



**Figure 87** Cache Flush Threshold Example

# Adding Capacity to the Disk Array

As your system storage requirements grow, you may need to increase the capacity of your disk array. Disk array capacity can be increased in any of the following ways:

- You can add new disk modules to the disk array if there are empty slots in the disk enclosures.
- You can add additional disk enclosures to the disk array.
- You can replace existing disk modules with higher capacity modules.

## Adding More Disk Modules

If there are empty slots in the disk enclosures, the easiest way to increase capacity is to add more disk modules and create additional LUNs. The disks you add will typically be used to bind LUNs, so make sure you add enough disk modules to create the desired structure. Disk modules can be added to the disk array while it is online without disrupting operation.

***CAUTION*** *Ensure that there is no important data on the disk module you are installing in the disk array. Any data on the disk will be lost when it is installed.*

To increase capacity by adding disk modules:

1. Install the disk modules in available slots in the disk enclosures. See "Disk Enclosure Modules" on page 386 for instructions on installing disk modules.

**Note** Each disk in the LUN must be in a different enclosure/channel to maintain high availability.

2. Bind a LUN with the new disks using the management tool of your choice,:
   - To use SAM, see "Binding a LUN" on page 267
   - To use Array Manager 60, see "Binding a LUN" on page 289
   - To use STM, see "Binding a LUN" on page 314

**Note** After binding a LUN, you must execute the `insf -e` command to install special device files on the LUN. This makes the LUN usable by the operating system.

3. Disable WCE on the disk modules using the following command:

   **amutil -w <ArrayID>**

   See "Disabling Disk Module Write Cache Enable (WCE)" on page 302 for more information.

## Moving Disks from One Disk Array to Another

***CAUTION*** *Before moving disks from one array to another, ensure that there is no important data on the disks. All data on the disks will be lost when they are installed in the new disk array.*

If you have more than one HP SureStore E Disk Array FC60, you can move disks from one array to another to balance capacity. The disks can be installed in the new array with power on and they will be treated as new disks. Any important data on the disks should be backed up before the disks are moved.

**Note** **Can I move a disk module within the array with power on?**
You should not move disks within the array with power on. Moving a disk from one slot to another will erase the data on the disk module.

## Adding Additional Disk Enclosures

Adding additional disk enclosures is another way to increase the capacity of the disk array. Up to six disk enclosures can be added to a disk array. Because it requires shutting down and possibly reconfiguring the disk array, adding new disk enclosures should be done by a trained-service representative. See "Adding Disk Enclosures to Increase Capacity" on page 224 for more information on adding disk enclosures.

An alternative to adding new disk enclosures is to move enclosures and their associated disks from one Disk Array FC60 to another. The LUNs on all disks in the enclosure should be unbound before moving it. See "Moving Disks from One Disk Array to Another" on page 255 for more details.

## Replacing Disk Modules with Higher Capacity Modules

If you have no available slots for adding new disk modules, you can increase the capacity of the disk array by replacing existing disk modules with higher capacity disk modules. This is typically done on a LUN-by-LUN basis and requires backing up and restoring all data on the LUN.

To increase capacity by replacing disk modules:

1. Identify the LUN impacted by the replacement of the disk modules. Make sure you replace all the disks in the LUN with the higher-capacity disks.
2. Backup all data on the LUN.
3. From the host, stop all I/O activity to the LUN unmount the file system.
4. Unbind the LUN using the management tool of your choice:
   - To use SAM, see "Unbinding a LUN" on page 271
   - To use Array Manager 60, see "Unbinding a LUN" on page 293
   - To use STM, see "Unbinding a LUN" on page 315
5. Replace the disk modules. See "Disk Enclosure Modules" on page 386 for instructions on replacing disk modules.

6. Bind a LUN with the new disks using the management tool of your choice:
   - To use SAM, see "Binding a LUN" on page 267
   - To use Array Manager 60, see "Binding a LUN" on page 289
   - To use STM, see "Binding a LUN" on page 314

**Note** After binding a LUN, you must execute the `insf -e` command to install special device files on the LUN. This makes the LUN usable by the operating system.

7. Mount the file system on the LUN and restore the data to the LUN.
8. Disable WCE on the disk modules using the following command:

   **amutil -w <ArrayID>**

   See "Disabling Disk Module Write Cache Enable (WCE)" on page 302 for more information.

# Upgrading Controller Cache to 512 Mbytes

Controller cache can be upgraded from the standard 256 Mbytes of cache to 512 Mbytes. This provides improved I/O performance for write operations. See Table 58 on page 416 for cache upgrade kit part numbers.

***CAUTION*** *The cache upgrade kit must be installed by service-trained personnel only. Attempting to install the cache upgrade kit without the proper training may damage the disk array controllers.*

### Determining What Type of Cache Memory is Installed on the Controllers

If there is currently 256 Mbytes of cache installed on the controllers, you must determine if this memory is configured as two 128 MB DIMMs, or as a single 256 MB DIMM. This will influence the number of upgrade kits you must order.

You can determine the cache configuration by checking the Manufacturing Part Number (MFG P.N.) on the product label on the front of the controller.

1. Remove the controller enclosure front panel.
2. Check the product label on the front of the controller. Depending on the part number, the cache is configured in one of the following ways:
   - MFG P.N. A5278-62006 uses two 128 MB DIMMs. It will be necessary to order two 256 DIMMs for each controller.
   - MFG P.N. A5278-62016 uses one 256 MB DIMM. It will be necessary to order one 256 DIMM for each controller.
3. Once you know the cache configuration, refer to Table 32 to determine how many cache upgrade kits are required.

**Table 32** Controller Cache Upgrade Kit Selection

| Initial controller configuration | Cache Upgrade Kit(s) |
|---|---|
| Dual controllers, each with two 128 MB DIMMs | Two A5279A kits |
| Dual controllers, each with one 256 MB DIMM | One A5279A kit |
| Single controller with two 128 MB DIMMs | One A5279A kit |
| Single controller with one 256 MB DIMM | One A5279A kit (only one of the DIMMs will be used) |

# Managing the Disk Array Using SAM

Most of the tasks involved in everyday management of the disk array can be performed using SAM. Using SAM you can:

- Check disk array status
- Bind and unbind LUNs
- Add and remove global hot spares

**Note**

**Does it make any difference which controller I select when managing the disk array?**

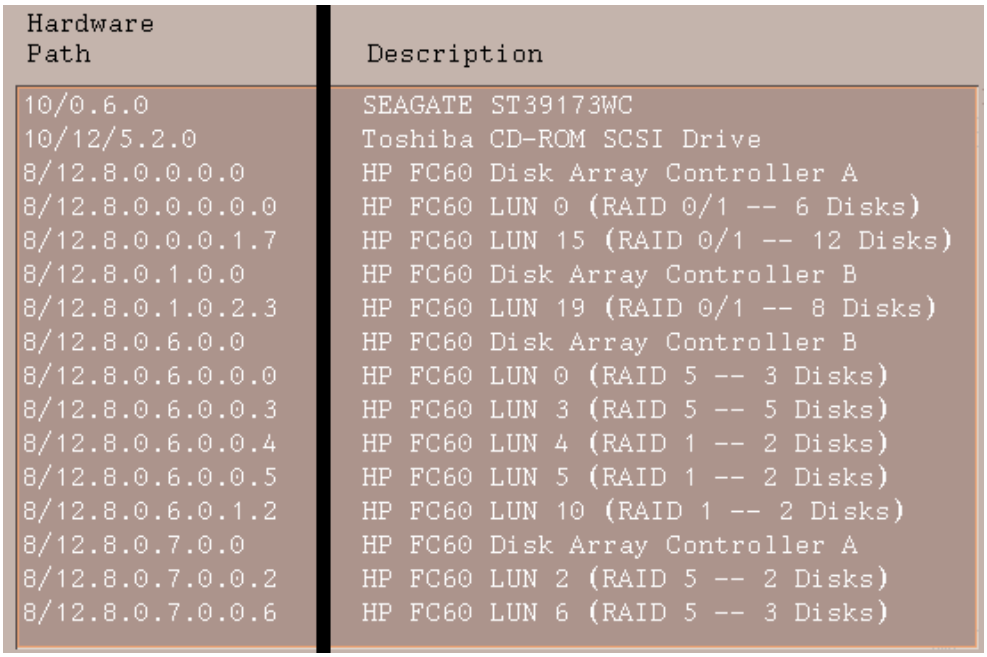
When using SAM, you must select one of the controllers on the disk array you want to manage. Typically you can select either controller to perform an operation and the outcome will be the same. An exception is when you are binding a LUN. The controller you select will be assigned ownership of the LUN. To maintain I/O performance, you should divide the ownership of the LUNs on the disk array between the two controllers.

## Checking Disk Array Status

All aspects of disk array operation are continually monitored and the current status is stored for viewing. You can selectively view the status of any portion of the disk array configuration.

To view disk array status:

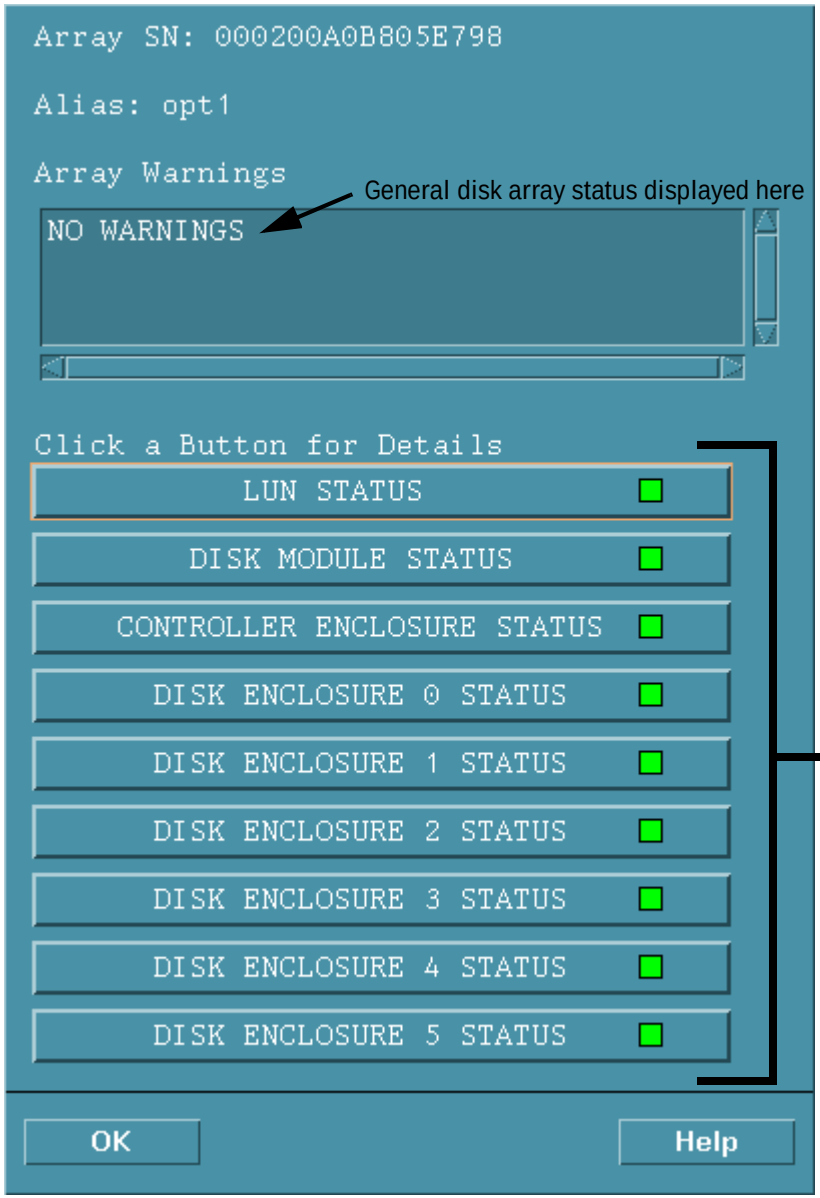
1. On the main SAM screen, double-click the Disks and File Systems icon.
2. On the Disks and File Systems screen, double-click the Disk Devices icon. The Disk Devices list is displayed. There is an entry for each HP FC60 disk array controller, and for each LUN on the disk array. The LUNs are listed beneath the owning controller.



| Hardware Path | Description |
|---|---|
| 10/0.6.0 | SEAGATE ST39173WC |
| 10/12/5.2.0 | Toshiba CD-ROM SCSI Drive |
| 8/12.8.0.0.0.0 | HP FC60 Disk Array Controller A |
| 8/12.8.0.0.0.0.0 | HP FC60 LUN 0 (RAID 0/1 -- 6 Disks) |
| 8/12.8.0.0.0.1.7 | HP FC60 LUN 15 (RAID 0/1 -- 12 Disks) |
| 8/12.8.0.1.0.0 | HP FC60 Disk Array Controller B |
| 8/12.8.0.1.0.2.3 | HP FC60 LUN 19 (RAID 0/1 -- 8 Disks) |
| 8/12.8.0.6.0.0 | HP FC60 Disk Array Controller B |
| 8/12.8.0.6.0.0.0 | HP FC60 LUN 0 (RAID 5 -- 3 Disks) |
| 8/12.8.0.6.0.0.3 | HP FC60 LUN 3 (RAID 5 -- 5 Disks) |
| 8/12.8.0.6.0.0.4 | HP FC60 LUN 4 (RAID 1 -- 2 Disks) |
| 8/12.8.0.6.0.0.5 | HP FC60 LUN 5 (RAID 1 -- 2 Disks) |
| 8/12.8.0.6.0.1.2 | HP FC60 LUN 10 (RAID 1 -- 2 Disks) |
| 8/12.8.0.7.0.0 | HP FC60 Disk Array Controller A |
| 8/12.8.0.7.0.0.2 | HP FC60 LUN 2 (RAID 5 -- 2 Disks) |
| 8/12.8.0.7.0.0.6 | HP FC60 LUN 6 (RAID 5 -- 3 Disks) |

**Note! The center portion of the screen has been removed for clarity.**

3. Select a controller for the appropriate disk array from the Disk Devices list.
4. Select the Actions menu, and the View More Information... menu option. The Main Status screen is displayed showing the overall status of the disk array.
5. Click the appropriate button to view the detailed status for the corresponding disk array component. If you need any help in interpreting the status information, access the online help.


Array SN: 000200A0B805E798
Alias: opt1
Array Warnings
General disk array status displayed here
NO WARNINGS
Click a Button for Details
LUN STATUS
DISK MODULE STATUS
CONTROLLER ENCLOSURE STATUS
DISK ENCLOSURE 0 STATUS
DISK ENCLOSURE 1 STATUS
DISK ENCLOSURE 2 STATUS
DISK ENCLOSURE 3 STATUS
DISK ENCLOSURE 4 STATUS
DISK ENCLOSURE 5 STATUS
OK
Help
Click here for detailed status of indicated component

### Interpreting Status Indicators

A common set of colored status indicators are used to convey the current operating status of each disk array component. The colors are interpreted as follows:

| | |
|---|---|
| Green | The component is functioning normally. On a disk it also indicates that the disk is part of a LUN. |
| Red | The component has failed or is not installed. |
| Blue | Used only for disks, indicates the disk is functioning normally and is a member of the hot spare disk group. |
| White | Used only for disks, indicates the disk is functioning normally and is a member of the unassigned disk group. |

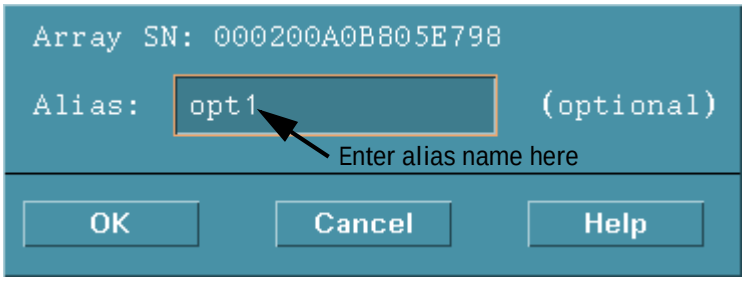
## Assigning an Alias to the Disk Array

If you have many disk arrays to manage, you may find it useful to assign an alias name to the disk arrays to help you in identifying them. The naming strategy you use may reflect the physical location of the disk array, or its function. For example, a disk array located in the data center might be assigned an alias of data_center_1.

To assign an alias:

1. On the main SAM screen, double-click the Disks and File Systems icon.
2. On the Disks and File Systems screen, double-click the Disk Devices icon. The Disk Devices list is displayed. There is an entry for each disk array controller.
3. Select a controller for the appropriate disk array from the Disk Devices list.

4. Select the Actions menu, the Disk Array Maintenance menu option, then Modify Array Alias. The Modify Array Alias screen is displayed.




5. Enter the name in the Alias field. An alias can contain up to 16 of the following characters: upper case letters, numbers, pound sign (#), period (.), and underscore (_). All other characters are invalid.
6. Click OK.

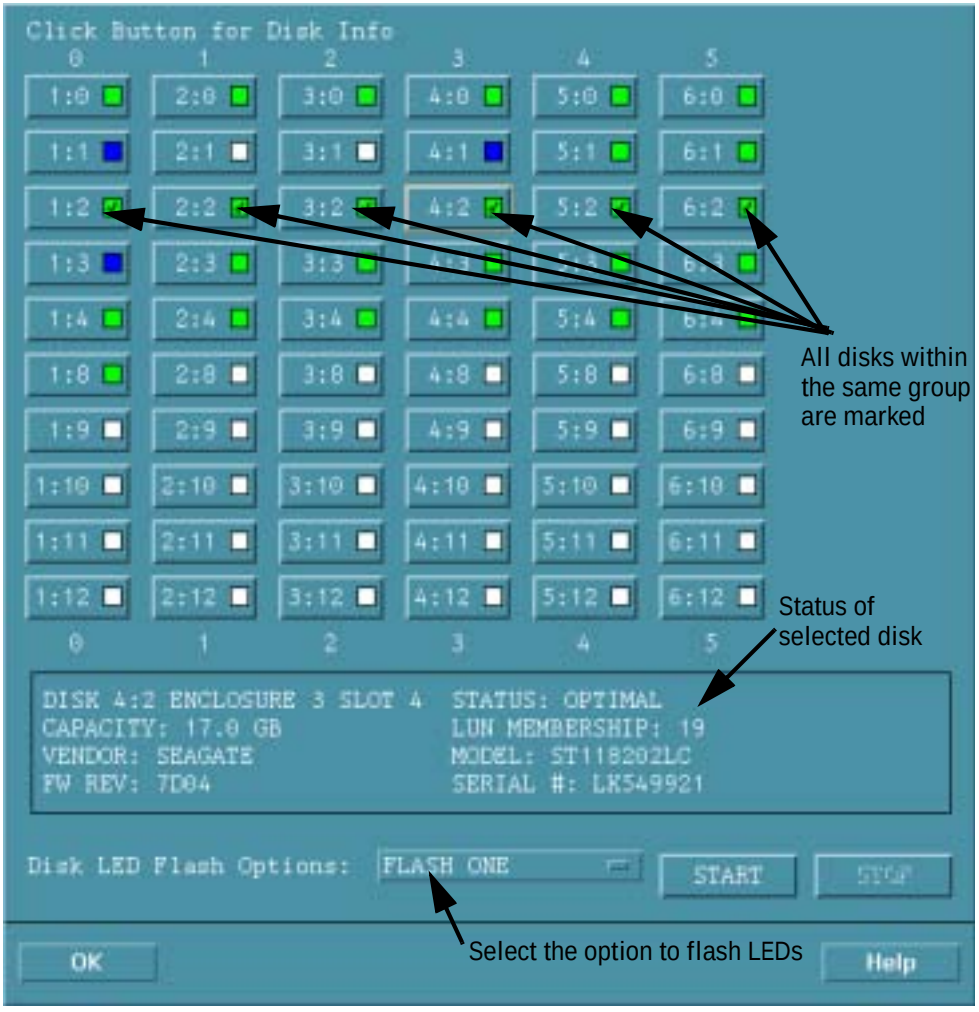
## Locating Disk Modules

To assist you in locating disk modules on the disk array, you can flash the amber Fault LED on one or multiple disks. This is useful if you want to identify the disks that make up a LUN, or a failed disk that must be replaced. It is advisable to positively identify a failed disk before removing it from the disk array. Removing the wrong disk could cause the LUN to become unavailable.

To locate a disk module:

1. On the main SAM screen, double-click the Disks and File Systems icon.
2. On the Disks and File Systems screen, double-click the Disk Devices icon. The Disk Devices list is displayed. There is an entry for each disk array controller.
3. Select a controller for the appropriate disk array from the Disk Devices list.
4. Select the Actions menu, and the View More Information... menu option. The Main Status screen is displayed.

5. Click the Disk Module Status button. The Disk Status window is displayed.


Click Button for Disk Info
0 1 2 3 4 5
1:0 2:0 3:0 4:0 5:0 6:0
1:1 2:1 3:1 4:1 5:1 6:1
1:2 2:2 3:2 4:2 5:2 6:2
1:3 2:3 3:3 4:3 5:3 6:3
1:4 2:4 3:4 4:4 5:4 6:4
1:8 2:8 3:8 4:8 5:8 6:8
1:9 2:9 3:9 4:9 5:9 6:9
1:10 2:10 3:10 4:10 5:10 6:10
1:11 2:11 3:11 4:11 5:11 6:11
1:12 2:12 3:12 4:12 5:12 6:12
0 1 2 3 4 5
All disks within the same group are marked
Status of selected disk
DISK 4:2 ENCLOSURE 3 SLOT 4 STATUS: OPTIMAL
CAPACITY: 17.0 GB LUN MEMBERSHIP: 19
VENDOR: SEAGATE MODEL: ST118202LC
FW REV: 7D04 SERIAL #: LK549921
Disk LED Flash Options: FLASH ONE
START
STOP
Select the option to flash LEDs
OK
Help

6. Select the disk you want to identify. A check mark will appear on the selected disk and all the other disks in the same disk group. For example, if the selected disk is part of a LUN, all disks within the LUN will be checked. If the disk is a global hot spare, all global hot spares will be checked.
7. Click on the Disk LED Flash Options button and select the desired option:
   - Flash One will flash the Fault LED on the selected disk only.
   - Flash Group will flash the Fault LEDs on all disk within the group
   - Flash Enclosure will flash the Fault LEDs on all disks in the same enclosure as the selected disk
   - Flash All flashes the Fault LEDs on all disks within the entire disk array
8. Click Start to begin flashing the Fault LEDs
9. Click Stop to stop flashing the Fault LEDs, or click OK to close the window and stop the flashing.

## Binding a LUN

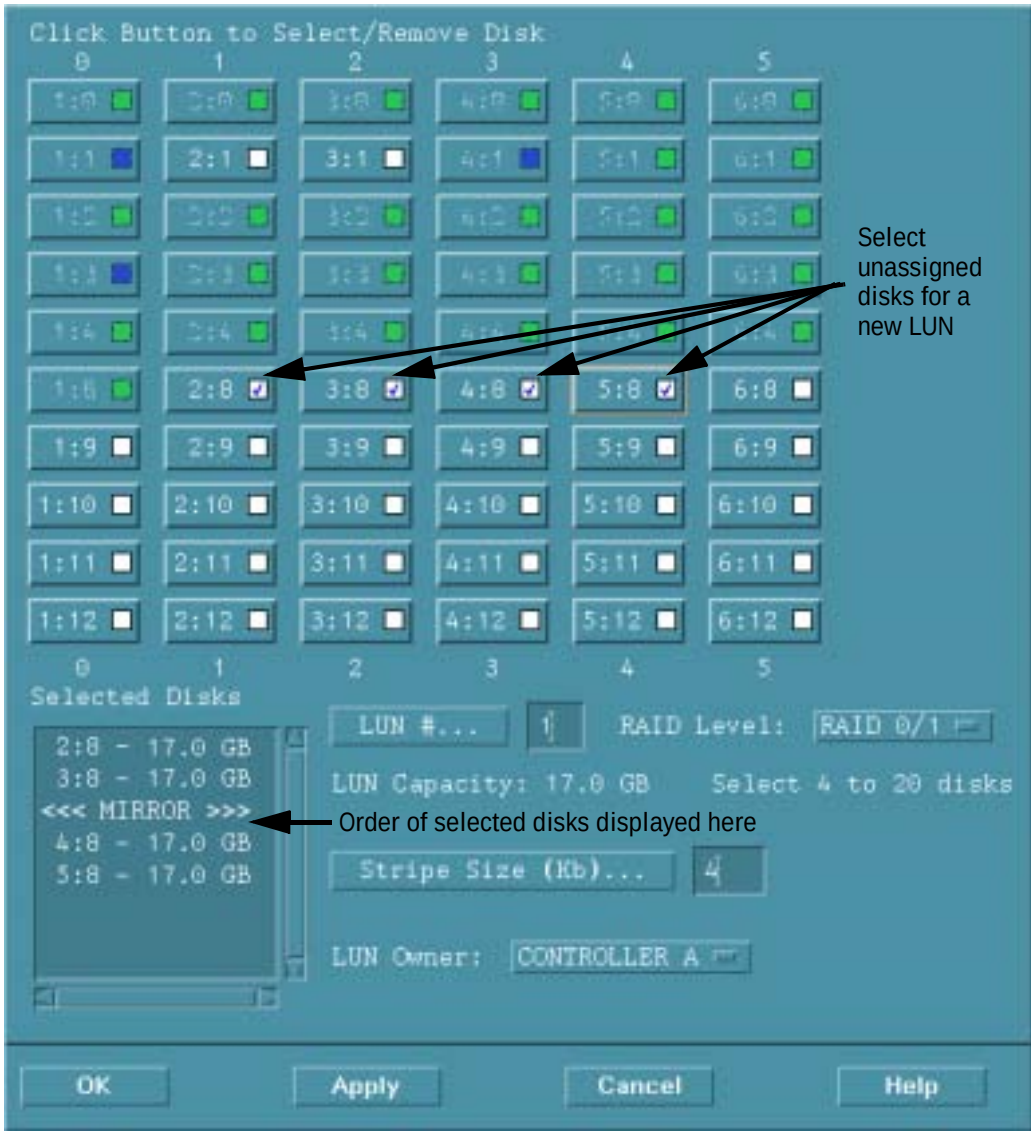
A number of settings allow you to define the LUN configuration. Before binding a LUN, make sure you understand what each of the settings does and how it will impact LUN behavior. See "Configuring LUNs" on page 242 for more information.

The LUN binding process is designed to maintain high-availability by recommending that you avoid creating a LUN that includes more than one disk per enclosure. If you attempt to do so, you will be alerted that the resulting LUN will be vulnerable to an enclosure failure.

**Note** **How long does it take to bind a LUN?**
The time it takes to bind a LUN depends on the size of the LUN you are creating. The larger the LUN, the longer it takes. For example, binding a RAID 5 LUN with 18-Gbyte disks can take up to 60 minutes.

To bind a LUN:

1. On the main SAM screen, double-click the Disks and File Systems icon.
2. On the Disks and File Systems screen, double-click the Disk Devices icon. The Disk Devices list is displayed. There is an entry for each disk array controller.
3. Select a controller for the appropriate disk array from the Disk Devices list.
4. Select the Actions menu, the Disk Array Maintenance menu option, and then Bind LUN. The LUN Management screen is displayed.


Click Button to Select/Remove Disk
0 1 2 3 4 5
1:8 2:8 3:8 4:8 5:8 6:8
1:9 2:9 3:9 4:9 5:9 6:9
1:10 2:10 3:10 4:10 5:10 6:10
1:11 2:11 3:11 4:11 5:11 6:11
1:12 2:12 3:12 4:12 5:12 6:12
Select unassigned disks for a new LUN
Selected Disks
2:8 - 17.0 GB
3:8 - 17.0 GB
<<< MIRROR >>>
4:8 - 17.0 GB
5:8 - 17.0 GB
Order of selected disks displayed here
LUN #...
1
RAID Level:
RAID 0/1
LUN Capacity: 17.0 GB
Select 4 to 20 disks
Stripe Size (Kb)...
4
LUN Owner:
CONTROLLER A
OK
Apply
Cancel
Help

5. Click the LUN # button and select the desired number for the LUN. You can also enter the LUN number directly in the field.
6. Click the RAID Level button and select the desired RAID level for the LUN.
7. Select the disks to include in the LUN. All unassigned disks are identified with a status of white.

   As you select disks, the capacity of the LUN is calculated and displayed, and the disks are added to the Selected Disks field. When creating a 0/1 LUN, the Selected Disks field identifies the data disks in the upper portion of the field and the associated mirror disks in the lower portion (below the <<<MIRROR>>> separator). See "Selecting Disks for a RAID 0/1 LUN" on page 243 for more information.
8. Click the Stripe Size button and select the stripe element size. You can also enter the stripe size in the field. The stripe size must be a multiple of the current cache page size. If you enter a value that is not a multiple of the cache page size, the value will be rounded up to the nearest correct value.
9. In the LUN Owner field, select the controller that will own the LUN. LUN ownership should be balanced between controllers for optimum performance.
10. Click OK to bind the LUN and exit the screen, or click Apply if you have additional LUNs to bind.

**Note** After binding a LUN, you must execute the `insf -e` command to install special device files for the LUN. This makes the LUN usable by the operating system.

## Unbinding a LUN

Unbinding a LUN makes its capacity available for the creation of a new LUN. All disks in the LUN are returned to the Unassigned disk group when the LUN is unbound.

***CAUTION*** *All data on a LUN is lost when it is unbound. Make sure you backup any important data on the LUN before unbinding it.*

*Unbinding a LUN will have the same effect on the host as removing a disk. Before unbinding a LUN, check your operating system documentation for any additional information or steps that may be required before unbinding a LUN.*

To unbind a LUN:

1. On the main SAM screen, select Disks and File Systems.
2. On the Disks and File Systems screen, double-click the Disk Devices icon. The Disk Devices list is displayed. There is an entry for each disk array controller, and each LUN on the disk array. The LUNs are listed beneath the owning controller.
3. From the Disk Devices list, select the LUN to be unbound.
4. Select the Actions menu, the Disk Array Maintenance menu option, and then Unbind.
5. Select "OK" to confirm the request.

## Replacing a LUN

It is possible to replace an existing LUN. This operation unbinds the LUN, and then rebinds it using any new parameters specified. This is useful if you wish to change the settings of an existing LUN, such as its size or RAID level, using a single command.

***CAUTION*** *All data on a LUN is lost when it is replaced. Make sure you backup any important data on the LUN before replacing it.*

**Note** **Can I replace any LUN on the disk array?**
Yes. In addition, the replace command is the only way you can alter the configuration of LUN 0.

LUN 0 is unique in that it must exist on the disk array to permit communication with the host. Consequently, you cannot unbind LUN 0. If you want to alter LUN 0, you must use the replace command.

To replace a LUN:

1. On the main SAM screen, select Disks and File Systems.
2. On the Disks and File Systems screen, double-click the Disk Devices icon. The Disk Devices list is displayed. There is an entry for each disk array controller, and each LUN on the disk array. The LUNs are listed beneath the owning controller.
3. From the Disk Devices list, select the LUN to be unbound.
4. Select the Actions menu, the Disk Array Maintenance menu option, and then Replace. The LUN Management screen is displayed.
5. Click the RAID Level button and select the desired RAID level for the LUN.
6. Select the disks to include in the LUN. The disks that were members of the original LUN are marked for inclusion. Remove any of these disks, or add additional disks as desired.
7. Click the Stripe Size button and select the stripe element size. You can also enter the stripe size in the field. The stripe size must be a multiple of the current cache page size. If you enter a value that is not a multiple of the cache page size, the value will be rounded up to the nearest correct value.
8. In the LUN Owner field, select the controller that will own the LUN. LUN ownership should be balanced between controllers for optimum performance.
9. Click OK to bind the new LUN and exit the screen.

**Note** After binding a LUN, you must execute the `insf -e` command to install special device files on the LUN. This makes the LUN usable by the operating system.

## Adding a Global Hot Spare

Global hot spares provide an additional level of protection for the data on your disk array. A global hot spare automatically replaces a failed disk, restoring redundancy and protecting against a second disk failure. For maximum protection against disk failures it is recommended that you add a global hot spare for each channel. For more information on using global hot spares, see "Global Hot Spares" on page 248.

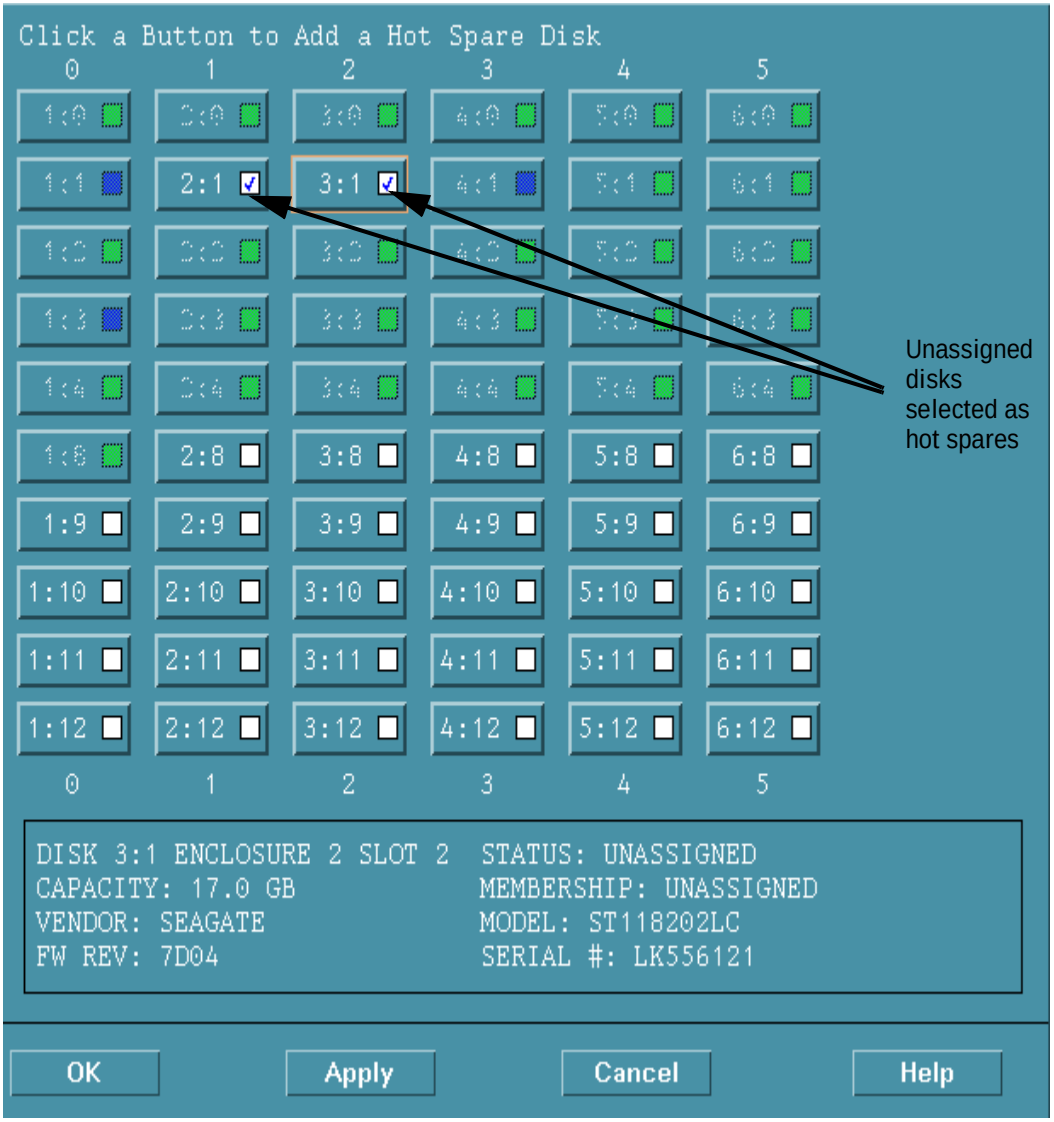
A global hot spare is added using an unassigned disk. If there are no unassigned disks available, you cannot add a global hot spare unless you install a new disk or unbind an existing LUN.

***CAUTION*** *If you have disks of different capacities in your disk array, always select disks of the largest capacity for your global hot spares. This ensures that any disk failure is protected. See "Global Hot Spare Disks" on page 61 for more information on selecting disks for global hot spares.*

*If a larger disk is used to replace a smaller disk that has failed, the difference in capacity will be unused. For example, if an 18-Gbyte disk is used as a global hot spare for a 9-Gbyte disk, the remaining 9 Gbytes of the global hot spare will be unused.*

To add a global hot spare:

1. On the main SAM screen, double-click the Disks and File Systems icon.
2. On the Disks and File Systems screen, double-click the Disk Devices icon. The Disk Devices list is displayed. There is an entry for each disk array controller.
3. Select a controller for the appropriate disk array from the Disk Devices list.
4. Select the Actions menu, the Disk Array Maintenance menu option, and then Add Hot Spare. The Add Hot Spare screen is displayed.


Click a Button to Add a Hot Spare Disk
0 1 2 3 4 5
2:1
3:1
1:12 2:12 3:12 4:12 5:12 6:12
DISK 3:1 ENCLOSURE 2 SLOT 2 STATUS: UNASSIGNED
CAPACITY: 17.0 GB MEMBERSHIP: UNASSIGNED
VENDOR: SEAGATE MODEL: ST118202LC
FW REV: 7D04 SERIAL #: LK556121
OK
Apply
Cancel
Help
Unassigned disks selected as hot spares

5. Select the disk to be used as a global hot spare.

   Only unassigned disks, identified by a white status indicator, are available for selection as hot spares.

6. Click OK to add the global hot spare and exit the screen, or click Apply if you want to add more global hot spares.

## Removing a Global Hot Spare

If you need to increase the available capacity of your disk array, you can do so by removing a global hot spare. The global hot spare will be added to the unassigned disk group and can be used to bind a LUN. However, the loss of the global hot spare may make your disk array more vulnerable when a disk failure occurs.

To remove a global hot spare:

1. On the main SAM screen, double-click the Disks and File Systems icon.
2. On the Disks and File Systems screen, double-click the Disk Devices icon. The Disk Devices list is displayed. There is an entry for each disk array controller.
3. Select a controller for the appropriate disk array from the Disk Devices list.
4. Select the Actions menu, the Disk Array Maintenance menu option, and then Remove Hot Spare. The Delete Hot Spare screen is displayed showing all current global hot spares, which are identified by a blue status indicator.
5. Select the global hot spare you want to remove.
6. Click OK to remove the global hot spare and exit the screen, or click Apply if you want to remove additional global hot spares.

# Managing the Disk Array Using Array Manager 60

The Array Manager 60 command line utilities allow you to configure, control, and monitor all aspects of disk array operation. Array Manager 60 is intended for performing the more advanced tasks involved in managing the disk array.

The Array Manager 60 utilities and the tasks they are used to perform are summarized in Table 33 and Table 34.

**Note** You must log in as superuser or root to use the Array Manager 60 utilities to manage the disk array.

**Table 33** Array Manager 60 Task Summary

| Task | Command |
|---|---|
| Capacity Management | |
| Binding a LUN | amcfg -L <cntrlrID>:<LUN> -d <channel:ID>,<channel:ID>..... -r <RAIDlevel> <options> <ArrayID> |
| Unbinding a LUN | amcfg -D <LUN> [-force] <ArrayID> |
| Calculating LUN Capacity | amcfg -C -d channel:ID,channel:ID... -r <RAIDlevel> -s <stripe> <ArrayID> |
| Adding a Global Hot Spare | ammgr -h channel:ID <ArrayID> |
| Removing a Global Hot Spare | ammgr -d channel:ID <ArrayID> |
| Identifying Disks | amdsp -d <ArrayID> |
| Changing LUN Ownership | amcfg -M <LUN> -c <cntrlrID> <ArrayID> |
| Replacing a LUN | amcfg -R <cntrlrID>:<LUN> -d <channel:ID>,<channel:ID>..... -r <RAIDlevel> <options> <ArrayID> |
| Disk Array Status | |
| Displaying Status Information | amdsp <option> <ArrayID> |
| Listing Disk Array IDs | amdsp -i |
| Rescanning for Disk Arrays | amdsp -R |

**Table 33** Array Manager 60 Task Summary (cont'd)

| Task | Command |
|---|---|
| Disk Array Configuration | |
| Assigning an Alias to the Disk Array | ammgr -D <ArrayAlias> <ArrayID> |
| Setting Cache Page Size | ammgr -p {4 \| 16} <ArrayID> |
| Setting the Cache Flush Threshold | ammgr -T <cntrlrID>:<percent> <ArrayID> |
| Setting the Cache Flush Limit | ammgr -L <cntrlrID>:<percent> <ArrayID> |
| Disabling Disk Module Write Cache Enable (WCE) | amutil -w <ArrayID> |
| Synchronizing the Controller Date and Time | ammgr -t <ArrayID> |
| Managing the Universal Transport Mechanism (UTM) | ammgr -u\|-U [-l <LUN>] <ArrayID> |
| Disk Array Maintenance | |
| Locating Disk Modules | amutil <option> <ArrayID> |
| Checking Rebuild Progress | amdsp -r <ArrayID> |
| Changing Rebuild Priority Settings | amutil -R <LUN> -f <Freq> -a <Amt> <ArrayID> |
| Performing a Parity Scan | ammgr -P <LUN> |
| Viewing Disk Array Logs | amlog [-s <StartTime>] [-e <EndTime>] [-d <LogDir>] [-a <ArrayID>] |
| Flushing Disk Array Log Contents | amutil -l <ArrayID> |
| Purging Controller Logs | amutil -p <ArrayID> |
| Resetting Battery Age | ammgr -b |

**Table 34** Array Manager 60 Command Summary

| Command | Tasks |
|---|---|
| amcfg | Binding a LUN<br>Unbinding a LUN<br>Changing LUN Ownership<br>Replacing a LUN<br>Calculating LUN Capacity |
| ammgr | Adding a Global Hot Spare<br>Removing a Global Hot Spare<br>Assigning an Alias to the Disk Array<br>Setting Cache Page Size<br>Setting the Cache Flush Threshold<br>Setting the Cache Flush Limit<br>Synchronizing the Controller Date and Time<br>Performing a Parity Scan<br>Resetting Battery Age<br>Managing the Universal Transport Mechanism (UTM) |
| amdsp | Identifying Disks<br>Displaying Status Information<br>Listing Disk Array IDs<br>Rescanning for Disk Arrays<br>Checking Rebuild Progress |
| amutil | Locating Disk Modules<br>Changing Rebuild Priority Settings<br>Disabling Disk Module Write Cache Enable (WCE)<br>Flushing Disk Array Log Contents |
| amlog | Viewing Disk Array Logs |

## Command Syntax Conventions

The following symbols are used in the command descriptions and examples in this chapter.

**Table 35** Syntax Conventions

| Symbol | Meaning |
|---|---|
| < > | Indicates a variable that must be entered by the user. |
| \| | Only one of the listed parameters can be used (exclusive OR). |
| [ ] | Values enclosed in these braces are optional. |
| { } | Values enclosed in these braces are required. |

## Array Manager 60 man pages

Online man pages are included for each Array Manager 60 command. The man page includes detailed information about the command and its usage.

**To access the man page for an Array Manager 60 command, type:**

**man <command_name>**

Substitute one of the Array Manager 60 utility names for **command_name**. For example, to access the amdsp man page, type:

**man amdsp**

## Quick Help

A quick listing of the syntax and available options for a command can be displayed by using the "?" option with the command. For example, for quick information about the ammgr utility, type:

**ammgr -?**

## Selecting a Disk Array and Its Components

When using Array Manager 60, you must select the disk array you will be managing. In addition, many commands also require you to identify the controller, disk, or LUN within the disk array that will be impacted by the command.

The command parameters used to select a disk array and its internal components are listed and described in Table 36.

**Note** **Does it make any difference which controller I select?**
There are two paths to the disk array — one for each controller. Typically you can select either controller to perform an operation and the outcome will be the same. An exception is when you are binding a LUN. The controller you select will be assigned ownership of the LUN. To maintain I/O performance, you should divide the ownership of the LUNs on the disk array between the two controllers.

**Table 36** Command Parameters for Identifying Disk Array Components

| Command parameter | Description |
|---|---|
| ArrayID | Identifies the disk array. The ArrayID may be either the disk array ID (its S/N), or an alias name assigned to the array. Each disk array must have a unique ArrayID. The disk array S/N and alias can be determined using the amdsp -i command. |
| cntrlrID | Identifies the controller within the disk array controller subsystem. Values can be either A or B, depending on which slot the controller is installed in. |
| channel:ID | Identifies a disk within the disk array. The channel value indicates the SCSI channel that the disk enclosure containing the disk is connected to. The ID value indicates the SCSI ID assigned to the slot in which the disk is installed.<br><br>Valid channel values are 1 - 6, and valid SCSI ID values are 0 - 4, 8 - 12. |
| LUN | Identifies a specific LUN on the disk array. Valid LUN values are 0 - 31. |

## Preparing to Manage the Disk Array

Before you begin using Array Manager 60 to manage your disk array for the first time, you may want to perform the following procedure. It will locate all the disk arrays on the host and allow you to assign an alias to each one. A short, meaningful alias should be easier to remember than the disk array ID when using the Array Manager 60 commands.

1. Rescan for all disk arrays on the host by typing:

   **amdsp -R**

2. List all the disk arrays by typing:

   **amdsp -i**

   Each disk array will be listed along with its ID.

3. Assign an alias to each disk array by typing:

   **ammgr -D <ArrayAlias> <ArrayID>**

   **ArrayAlias** is the name you want to assign to the disk array. You can use up to 16 of the following characters: letters, numbers, pound sign (#), period (.), and underscore (_). All other characters are invalid.

## Checking Disk Array Status

An important part of managing the disk array involves monitoring its status to ensure it is working properly. Changes in disk array status may indicate a possible hardware failure, so it is important to check disk array status regularly.

All aspects of disk array operation are continually monitored and the current status is stored for viewing. You can selectively view the status of any portion of the disk array configuration.

See "Status Conditions and Sense Code Information" on page 317 for assistance in interpreting status information.

### Displaying Status Information

To display status information, type:

**amdsp <option> <ArrayID>**

- **option** identifies the type of status information to display**.** Table 37 lists the command options and the status information each displays.

**Table 37** Command Options for Displaying Disk Array Status

| Option | Status Information Displayed |
|---|---|
| none | General array information. This includes disk capacity usage. Sample output is shown in Figure 88. |
| -l [LUN] | LUN status. Status for all LUNs on the disk array will be displayed, unless a specific LUN is specified using the LUN option. If UTM is enabled, it will be displayed in the LUN output. |
| -g | Disk group status. |
| -d | Disk status. Complete information will be displayed for each disk in the disk array. |
| -c | Controller status. Information is displayed for both disk array controllers and for all disk system BCC modules. |
| -s | Array status. Information about the entire array is displayed. |
| -h | Hardware status. Information is displayed for all disk array hardware components, including fans and power supplies. |

**Table 37** Command Options for Displaying Disk Array Status (cont'd)

| Option | Status Information Displayed |
|---|---|
| -a | All status. This option displays all the information returned by the preceding options. |
| -p <devicefile> | Hardware path information. Displays hardware path information for the controller corresponding to the specified device file. |
| -r | Rebuild status. Display the progress of any rebuilds occurring on the disk array. |
| -A | Array server status. |

### Command Example

The following example all status information for disk array Array1.

**amdsp -a Array1**

A sample out put from this command is shown in Figure 88. The important fields are identified with the normal values.

**Note**

**Missing UTM LUN 31**
If the required patches for HP08 firmware are not installed on the host, the host will not recognize or display UTM LUN 31. If LUN 31 is not displayed, install the required patches as described in "Installing the Disk Array FC60 Software" on page 214.

**Figure 88** Disk Array Sample Status Output (amdsp)

```
Vendor ID           = HP
Product ID          = A5277A
Array ID            = 000A00A0B80673A6
Array alias         = Array1
----------------------------------------
Array State           = READY
Server name           = speedy
Array type            = 3
Mfg. Product Code     = 348-0040789

--- Disk space usage --------------------
Total physical               = 271.4 GB
Allocated to LUNs            = 135.4 GB
Used as Hot spare            = 0.0 GB
Unallocated (avail for LUNs) = 0.0 GB
----------------------------------------

Vendor ID           = HP
Product ID          = A5277A
Array ID            = 000A00A0B80673A6
Array alias         = Array1
----------------------------------------
LUN       Status        Capacity   Ctrl  RAID  Segment  Disks
---  -----------------  ---------  ----  ----  -------  -----
  0  OPTIMAL              16.9 GB   A     1          4  1:0
                                                       3:0
  1  OPTIMAL              16.9 GB   B     1          4  2:0
                                                       4:0
   .
   .
 31  UTM:GOOD

   LUN  WCE  RCD  CME  CWOB      WCA  RCA  CMA
   ---  ---  ---  ---  ----      ---  ---  ---
     0   X         X             X    X    X
     1   X         X             X    X    X
   .
   .
   .
Total capacity of LUNs on controller A = 84.6 GB
Total capacity of LUNs on controller B = 50.8 GB
Total capacity of all configured LUNs = 135.4 GB
```

LUN Information - the status of each LUN should be OPTIMAL

LUN 31 is the default UTM LUN. GOOD status indicates the UTM is enabled and operating normally. See note on page 283.

Cache Status for each LUN - all LUNs should have the status values shown

LUNs should typically be distributed between both controllers for best performance. LUN ownership can be reassigned to balance LUN usage.

```
Vendor ID           = HP
Product ID          = A5277A
Array ID            = 000A00A0B80673A6
Array alias         = Array1
----------------------------------------

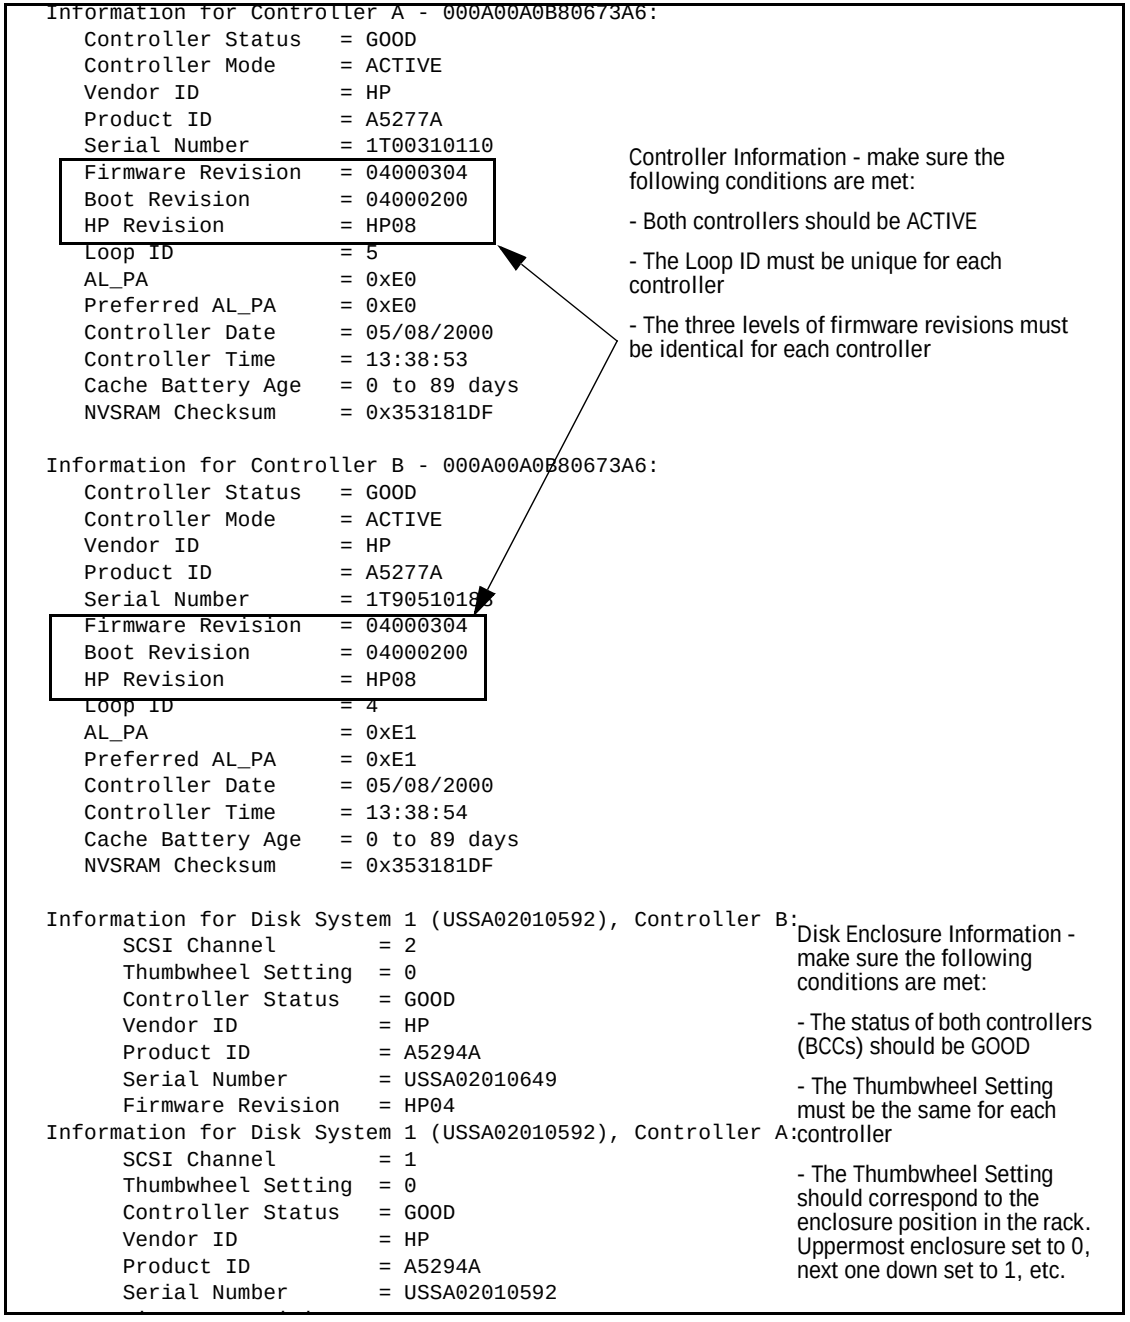
SCSI Channel:ID        = 1:0
Enclosure              = 0
Slot                   = 0
Disk State             = OPTIMAL
Disk Group and Type    = 060E86000238C6360F LUN
Capacity               = 17.0 GB
Manufacturer and Model = SEAGATE ST318203LC
Serial Number          = LRB61150
Firmware Revision      = HP01
   .
   .
   .
Total capacity of all installed physical disks = 271.4 GB

Hot Spare Activity
--------------------
None

Vendor ID           = HP
Product ID          = A5277A
Array ID            = 000A00A0B80673A6
Array alias         = Array1
----------------------------------------
Disk Group 060E86000738C6395B:
   Disk Group Type    = LUN
   Number of LUNs     = 1
   LUN(s)             = 6
   Remaining Capacity = 0.0 MB
   RAID Level         = 1
   Segment Size       = 4 KB
   Disks:
   1:3  3:3
   .
   .
   .
Total remaining capacity for LUN disk groups = 0.0 MB
```

Disk Information - the status of each disk should be Optimal for each disk assigned to a LUN.

Disk Group Information - this is another representation of the LUN information


Information for Controller A - 000A00A0B80673A6:
Controller Status = GOOD
Controller Mode = ACTIVE
Vendor ID = HP
Product ID = A5277A
Serial Number = 1T00310110
Firmware Revision = 04000304
Boot Revision = 04000200
HP Revision = HP08
Loop ID = 5
AL_PA = 0xE0
Preferred AL_PA = 0xE0
Controller Date = 05/08/2000
Controller Time = 13:38:53
Cache Battery Age = 0 to 89 days
NVSRAM Checksum = 0x353181DF
Information for Controller B - 000A00A0B80673A6:
Controller Status = GOOD
Controller Mode = ACTIVE
Vendor ID = HP
Product ID = A5277A
Serial Number = 1T90510183
Firmware Revision = 04000304
Boot Revision = 04000200
HP Revision = HP08
Loop ID = 4
AL_PA = 0xE1
Preferred AL_PA = 0xE1
Controller Date = 05/08/2000
Controller Time = 13:38:54
Cache Battery Age = 0 to 89 days
NVSRAM Checksum = 0x353181DF
Information for Disk System 1 (USSA02010592), Controller B:
SCSI Channel = 2
Thumbwheel Setting = 0
Controller Status = GOOD
Vendor ID = HP
Product ID = A5294A
Serial Number = USSA02010649
Firmware Revision = HP04
Information for Disk System 1 (USSA02010592), Controller A:
SCSI Channel = 1
Thumbwheel Setting = 0
Controller Status = GOOD
Vendor ID = HP
Product ID = A5294A
Serial Number = USSA02010592
Controller Information - make sure the following conditions are met:
- Both controllers should be ACTIVE
- The Loop ID must be unique for each controller
- The three levels of firmware revisions must be identical for each controller
Disk Enclosure Information - make sure the following conditions are met:
- The status of both controllers (BCCs) should be GOOD
- The Thumbwheel Setting must be the same for each controller
- The Thumbwheel Setting should correspond to the enclosure position in the rack. Uppermost enclosure set to 0, next one down set to 1, etc.

```
Vendor ID           = HP
Product ID          = A5277A
Array ID            = 000A00A0B80673A6
Array alias         = Array1
----------------------------------------
Overall State of Array          = READY
Array configuration:
   Cache Block Size             = 4 KB / 4 KB
   Cache Flush Threshold        = 80 % / 80 %
   Cache Flush Limit            = 100 % / 100 %
   Cache Size                   = 256MB / 256MB

Vendor ID           = HP
Product ID          = A5277A
Array ID            = 000A00A0B80673A6
Array alias         = Array1
----------------------------------------
Array Controller Subsystem:
    Controller A: GOOD
    Controller B: GOOD
            PS 1: GOOD
            PS 2: GOOD
           Fan 1: GOOD
           Fan 2: GOOD
     Temp Sensor: GOOD
         Battery: GOOD
Disk System 1 - USSA02010592:
BCC Controller B: GOOD
BCC Controller A: GOOD
            PS 1: GOOD
            PS 2: GOOD
           Fan 1: GOOD
           Fan 2: GOOD
     Temp Sensor: GOOD
Disk System 2 - USSA02010595:
BCC Controller B: GOOD
BCC Controller A: GOOD
            PS 1: GOOD
            PS 2: GOOD
           Fan 1: GOOD
           Fan 2: GOOD
     Temp Sensor: GOOD
No LUNs are currently rebuilding
```

Cache settings

Component status - all hardware components should be GOOD

## Listing Disk Array IDs

You may find it useful to list the disk arrays recognized by the host. The list will include both the disk array ID (or S/N) and alias name if one has been assigned. This is a quick way to determine the ID of each disk array on your system.

**Note** **What if a disk array is not listed?**
If the list does not reflect the current disk arrays on your system, rescan for disk array as described in the next procedure. This will update the disk array list.

To list all the disk arrays, type:

**amdsp -i**

**Note** To avoid having to use the array ID when managing the disk array, you may want to create an alias. You can also export the array ID to a variable and use the variable. For example, to export the ID for array `000000057D22` to a variable named id, type the following:

```
# export <id>=000000057D22
```

You can then use `$id` in place of `<array-id>`

## Rescanning for Disk Arrays

If the configuration of the disk arrays on your system has changed, you can rescan the disk arrays to update the information. This command causes Array Manager 60 to rescan the system for new or modified disk arrays connected to the host.

To rescan for disk arrays, type:

**amdsp -R**

When the rescan is complete, the command prompt will return. The command does not return a message indicating the operation is complete.

## Managing LUNs

Using Array Manager 60 you can perform the following tasks:

- Binding and unbinding LUNs
- Calculating LUN capacity
- Changing LUN ownership
- Replacing a LUN

### Binding a LUN

Binding LUNs is one of the most common tasks you will perform in managing the disk array. A number of settings allow you to define the LUN configuration. Before binding a LUN, make sure you understand what each of the settings does and how it will impact LUN behavior. See "Configuring LUNs" on page 242 for factors that may influence how you bind a LUN.

Before binding a LUN, you may want to calculate how much capacity the LUN will provide with the disks and settings you intend to use. Refer to "Calculating LUN Capacity" on page 292 for instructions on performing this task.

**Note** **How long does it take to bind a LUN?**
The time it takes to bind a LUN depends on the size of the LUN you are creating. The larger the LUN, the longer it takes. For example, binding a RAID 5 LUN with 18-Gbyte disks can take up to 60 minutes.

To bind a LUN, type:

**amcfg -L <cntrlrID>:<LUN> -d <channel:ID>,<channel:ID>,<channel:ID>..... -r <RAIDlevel> <options> <ArrayID>**

- **cntrlrID** - identifies the controller (A or B) that will assume ownership of the LUN.
- **LUN** - number assigned to the LUN. It must not conflict with an existing LUN number.
- **channel:ID** - identifies the channel and SCSI ID of each disk included in the LUN. Only unassigned disks can be used for the LUN. Multiple disks must be specified with no spaces between each disk. See "Selecting Disks for a RAID 0/1 LUN" on page 243 for more information if you are binding a RAID 0/1 LUN.

- **RAIDlevel** - RAID level used for the LUN. Valid RAID levels are 0, 1, and 5. RAID 0 support requires firmware version HP08 or later. A RAID 0/1 LUN is created by selecting RAID 1 with more than two disks.
- **<options> -** options giving you control over how the LUN is configured. Table 38 lists valid options and what they do.

**Note** After binding a LUN, you must execute the `insf -e` command to install special device files for the LUN. This makes the LUN usable by the operating system.

**Table 38** Command Options for Binding a LUN

| Option | Description | Default |
|---|---|---|
| -c capacity | Specifies the capacity of the LUN. Capacity can be specified in megabytes (M) or gigabytes (G) by appending the appropriate letter to the capacity value.<br><br>By default, the LUN will use the entire capacity available from all the disks. It is possible to specify a lower value for LUN capacity, but this will result in unused disk capacity. Any capacity not included in the LUN will be inaccessible and is essentially wasted capacity. | If not specified, the LUN capacity will equal the total available capacity of the specified disks.<br><br>If a capacity value is specified but no letter is appended, gigabytes will be assumed. |

**Table 38** Command Options for Binding a LUN (cont'd)

| Option | Description | Default |
|---|---|---|
| -force | Allows you to bind a LUN using two or more disks in the same enclosure or on the same channel. This option allows you to override the high-availability protection designed into the LUN binding process. Using this option you can specify more than one disk per enclosure or channel. You can also use this option to create a RAID 5 LUN that includes more than six disks. | If not specified, you cannot bind a LUN using multiple disks in the same enclosure. You also cannot bind a RAID 5 LUN with more than six disks. |
| -s segmentsize | Indicates the stripe segment size in Kbytes. The stripe segment size must be a multiple of the current cache page size setting. For example, if the cache page size is set to 4 Kbytes (8 blocks), valid settings would be 4, 8, 16, 32, 64, 128, etc. | If not specified, the stripe segment size defaults to the same block size used for the cache pages (4 Kbytes or 16 Kbytes). The default cache page size is 4 Kbytes. |

### Command Examples

The following example binds a 5-disk, RAID 5 LUN on disk array 0000005EBD20. The disk array includes five disk enclosures, each on its own channel. The LUN uses a stripe segment size of 16 Kbytes, is assigned LUN number 2, and is owned by controller A. Note that each disk is in a different enclosure for high availability, and that there are no spaces between the individual disk parameters.

**amcfg -L A:2 -d 1:1,2:2,3:1,4:4,5:3 -r 5 -s 16 0000005EBD20**

The following example binds a RAID 0/1 LUN on disk array rack_51. The LUN is owned by controller B, is assigned LUN number 4, and uses a stripe segment size of 4 Kbytes. Note that the disks selected create mirrored pairs that are in separate enclosures (1:3, 2:3 and 1:4, 2:4). This maintains high availability. Although RAID 1 is specified, the inclusion of more than two disks causes the disk array to create a RAID 0/1 LUN.

**amcfg -L B:4 -d 1:3,1:4,2:3,2:4 -r 1 -s 4 rack_51**

## Identifying Disks

Binding a LUN requires the use of unassigned disks. If you are not sure which disks are unassigned, you can determine which disks are available.

To identify unassigned disks, type

**amdsp -d <ArrayID>**

The status of all disks in the array will be returned. The information includes the disk group the disk is a member of. Disks in the Unassigned disk group can be used for binding a LUN. You may want to print the information or write down the unassigned disks before you begin binding the LUN.

## Calculating LUN Capacity

Before actually binding a LUN, you may want to calculate how much capacity will be available for a given set of configuration values. This command does not bind a LUN; it simply returns the capacity that would be available if you had bound a LUN using the values provided. You can alter the settings to see how LUN capacity is affected until you get the desired results. Using this command before you bind a LUN will ensure that you achieve the correct capacity.

To calculate the capacity of a LUN, type:

**amcfg -C -d channel:ID,channel:ID... -r <RAIDlevel>**
**-s <stripe> <ArrayID>**

- **RAIDlevel** - the RAID level used for the LUN. Valid RAID levels are 1, 0/1, and 5.
- **stripe** - the size of the stripe element in Kbytes.

### Command Example

The following example calculates the capacity of a RAID 5 LUN using four disks and a stripe element size of 16 Kbytes.

**amcfg -C -d 1:1,2:3,3:1,4:2 -r 5 -s 16 0000005EBD20**

## Unbinding a LUN

Unbinding a LUN makes its capacity available for the creation of a new LUN. All disks assigned to the LUN are returned to the Unassigned disk group when the LUN is unbound.

***CAUTION*** *All data on a LUN is lost when it is unbound. Make sure you backup any important data on the LUN before unbinding it.*

*Unbinding a LUN will have the same impact on the host as removing a disk. Before unbinding a LUN, check your operating system documentation for any additional information or steps that may be required after unbinding a LUN.*

To unbind a LUN, type:

**amcfg -D <LUN> [-force] <ArrayID>**

- **LUN** - number assigned to the LUN you want to unbind
- The **-force** option is required to delete LUN 0. LUN 0 may only be deleted if a UTM LUN has been enabled via NVSRAM setting. If UTM is disabled, LUN 0 is required for the SCSI command path. The UTM LUN itself may not be deleted using this command.

### Command Example

The following example unbinds LUN 5 on disk array rack_1.

**amcfg -D 5 rack_1**

## Changing LUN Ownership

When binding a LUN, one of the disk array controllers is identified as the owner of the LUN. It is possible to change the ownership of a bound LUN dynamically - that is, without unbinding and rebinding the LUN. You may want to change the ownership of a LUN if you notice an imbalance of I/O loading between controllers. Controller ownership will change automatically if a failure occurs in the primary I/O path to the LUN.

**Note** **Does the primary path selected using LVM impact LUN ownership?**
Yes. The primary path established using LVM defines the owning controller for the LUN. This may override the controller ownership defined when the LUN was bound. For example, if controller A was identified as the owning controller when the LUN was bound, and LVM subsequently established the primary path to the LUN through controller B, controller B becomes the owning controller.

To change the controller ownership of a LUN, type:

**amcfg -M <LUN> -c <cntrlrID> <ArrayID>**

- **cntrlrID** identifies the controller (A or B) that is being assigned LUN ownership

#### Command Example

The following example changes the ownership of LUN 3 to controller B on disk array 0000005EBD20.

**amcfg -M 3 -c B 0000005EBD20**

### Replacing a LUN

It is possible to replace an existing LUN. This operation unbinds the LUN and then rebinds it using the new parameters specified. This is useful if you wish to change the settings of an existing LUN, such as its size or RAID level, using a single command.

***CAUTION*** *All data on a LUN is lost when it is replaced. Make sure you backup any important data on the LUN before replacing it.*

**Note** **Can I replace any LUN on the disk array?**
Although you can replace any LUN, the replace command is intended to be used with LUN 0. LUN 0 is unique in that it must exist on the disk array to permit communication with the host. Consequently, you cannot unbind LUN 0. If you want to alter LUN 0, you must use the replace command. For all other LUNs, it is recommended that you use the unbind and bind commands to alter LUN configuration.

To replace a LUN, type:

**amcfg -R <cntrlrID>:<LUN> -d <channel:ID>,<channel:ID>.....**
**-r <RAIDlevel> <options> <ArrayID>**

- The parameters and options available when replacing a LUN are the same as those used when binding a LUN. See "Binding a LUN" on page 289.

### Command Examples

The following example replaces existing LUN 0 on disk array 0000005EBD20. The new LUN is RAID 5, uses a stripe segment size of 16 Kbytes, and is owned by controller A. Note that each disk is in a different enclosure for high availability, and that there are no spaces between the individual disk parameters.

**amcfg -R A:0 -d 1:1,2:2,3:1,4:4,5:3 -r 5 -s 16 0000005EBD20**

## Managing Global Hot Spares

Global hot spares provide an additional level of protection for the data on your disk array. A global hot spare automatically replaces a failed disk, restoring data redundancy and performance that may suffer while in degraded mode. For maximum protection against disk failure it is recommended that you add one global hot spare per channel.

For more information on using global hot spares, see "Global Hot Spares" on page 248.

***CAUTION*** *If you have disks of different capacities in your disk array, always select disks of the largest capacity for your global hot spares. This ensures that any disk failure is protected. See "Global Hot Spare Disks" on page 61 for more information on selecting disks for global hot spares.*

*If a larger disk is used to replace a smaller disk that has failed, the difference in capacity will be unused. For example, if an 18-Gbyte disk is used as a global hot spare for a 9-Gbyte disk, the remaining 9 GB of the global hot spare remain unused.*

## Adding a Global Hot Spare

A global hot spare is added using an unassigned disk. If there are no unassigned disks available, you cannot add a global hot spare unless you install a new disk or unbind an existing LUN.

To add a global hot spare, type:

**ammgr -h channel:ID <ArrayID>**

### Command Example

The following example adds a global hot spare using disk 1:1 on disk array 0000005EBD20.

**ammgr -h 1:1 0000005EBD20**

## Removing a Global Hot Spare

If you need to increase the available capacity of your disk array, you can do so by removing a global hot spare. The global hot spare will be added to the unassigned disk group and can be used to bind a LUN. However, the loss of the global hot spare may make your disk array more vulnerable when a disk failure occurs.

To remove a global hot spare, type:

**ammgr -d channel:ID <ArrayID>**

### Command Example

The following example removes disk 2:3 as a global hot spare on disk array rack_51.

**ammgr -d 2:3 rack_51**

## Managing Disk Array Configuration

### Assigning an Alias to the Disk Array

If you have many disk arrays to manage, you may find it useful to assign an alias name to each disk array to help you identify them. A short, meaningful alias should be easier to remember than the disk array ID when using the Array Manager 60 commands. The naming strategy you use may reflect the physical location of the disk array or its function. Once you have assigned an alias to the disk array, you can use the new name in the <ArrayID> field used to identify the disk array when using a command.

To assign an alias to the disk array, type:

**ammgr -D <ArrayAlias> <ArrayID>**

- **ArrayAlias** is the name you want to assign to the disk array. An alias can contain up to 16 of the following characters: letters, numbers, pound sign (#), period (.), and underscore (_). All other characters are invalid.

#### Command Example

The following example assigns an alias of rack_1 to disk array 0000005EBD20:

**ammgr -D rack_1 0000005EBD20**

### Synchronizing the Controller Date and Time

During installation, the date and time maintained internally by the disk array should have been synchronized to the host. This ensures that any information created by the disk array such as log entries reflect the proper time they occurred. If necessary, the disk array date and time can be resynchronized with the host.

To synchronize the controller date and time with the host, type:

**ammgr -t <ArrayID>**

## Managing the Universal Transport Mechanism (UTM)

On firmware HP08 and later, the Universal Transport Mechanism (UTM) serves as the SCSI communication path between the host and the disk array. In earlier versions of firmware, this communication was done using LUN 0. The UTM is configured as a separate LUN, which is used only for communication and not for storing data. Because it consumes one of the available LUNs, only 31 LUNs are available when using the UTM. Any attempt to perform an I/O to the UTM LUN will be rejected.

The UTM is enabled by default and configured as LUN 31. It is possible to change the LUN used for the UTM to something other than 31. When upgrading to HP08 firmware on a disk array that already has LUN 31 defined, it will be necessary to copy the data a new LUN and delete LUN 31 before downloading the HP08 firmware.

One of the primary benefits of using the UTM is Major Event Logging (MEL). If the UTM is disabled, major event logging will not be performed. Instead, the standard controller logs will be used to gather disk array status. Because of the benefits it offers, it is recommended that the UTM be enabled.

### Enabling and Changing the UTM LUN

LUN 31 is assigned as the default LUN used for the UTM. In most cases this should be acceptable. It is possible to change the LUN used for the UTM if you wish to do so. This command is also used to enable the UTM if it is disabled.

**Note** LUN 0 cannot be used for the UTM.

To change the LUN number used for UTM, or to enable the UTM, type:

**ammgr -U -l <LUN> <ArrayID>**

- **LUN** identifies the default LUN that will be used for the ATM. This value is stored in controller NVSRAM. If **LUN** is not specified, the default value currently stored in NVSRAM will be used for the ATM. The UTM will be enabled if it is currently disabled.

**Note** After executing the above command, the disk array controllers must be manually reset or power cycled before the new setting will be invoked. When the power on completes, execute the following commands:

```
ioscan
insf -e
amdsp -R
```

### Disabling the UTM

Although it possible to disable the UTM, it is not recommended that you do so. The benefits provided by the UTM, such as major event logging, are not realized when the UTM is disabled. If you want to disable major event logging, you can do so by disable the UTM.

**Note** It is necessary to disable the UTM LUN before converting from firmware HP08 to ealier versions of firmware.

To disable use of the UTM, type:

**ammgr -u <ArrayID>**

**Note** After executing the above command, the disk array controllers must be manually reset or power cycled before the new setting will be invoked. When the power on completes, execute the following commands:

```
ioscan
insf -e
amdsp -R
```



## Managing Cache Parameters

Several cache parameters control how the disk array manages controller cache memory. These setting have a direct influence on I/O performance. Before altering these settings,

see Table 31 on page 250 for details on what performance impact altering these settings may have.

## Setting Cache Page Size

Data caching improves disk array performance by storing data temporarily in cache memory. The cache page size can be set to either 4 Kbytes or 16 Kbytes. The default page size is 4 Kbytes. The page size is set for both controllers in the disk array.

To set cache page size, type:

**ammgr -p {4|16} <ArrayID>**

### Command Example

The following example sets the cache page size to 16 Kbytes on disk array rack_1:

**ammgr -p 16 rack_1**

## Setting the Cache Flush Threshold

The cache flush threshold sets the level at which the disk array begins flushing (writing) cache content to the disks. Cache will be flushed until the cache flush limit value is reached. The default value for this setting is 80%. The cache flush threshold can be set independently for each controller.

To set cache flush threshold, type:

**ammgr -T <cntrlrID>:<percent> <ArrayID>**

- **percent** indicates the percentage of cache memory that must be full before flushing begins. For example, if set to 80%, cache flushing will begin when 80% of cache memory is full.

### Command Example

The following example sets the cache flush threshold to 80% for controller A on disk array 0000005EBD20:

**ammgr -T A:80 0000005EBD20**

### Setting the Cache Flush Limit

Sets the amount of unwritten data to remain in cache after a flush is completed on the given controller. The cache flush limit sets the level at which the disk array stops flushing cache contents to the disks. This value is expressed as a percentage of the current cache flush threshold. The default value for this setting is 100%. The cache flush limit can be set independently for each controller.

To set cache flush limit, type:

**ammgr -L <cntrlrID>:<percent> <ArrayID>**

- **percent** indicates the percentage of cache memory that remains full when flushing of cache contents stops. For example, if set to 90%, cache memory will be flushed until 90% of the cache amount defined by the cache flush threshold remains full.

#### Command Example

The following example sets the cache flush limit to 100% for controller A on disk array 0000005EBD20:

**ammgr -L A:100 0000005EBD20**

## Disabling Disk Module Write Cache Enable (WCE)

**Note** To ensure optimum protection against data loss, it is recommended that Write Cache Enable be disabled on all disks in the array. Disabling disk WCE will impact disk array performance. However, it reduces the potential for data loss during a power loss.

To ensure optimum data integrity, it is recommended that the Write Cache Enable (WCE) feature be disabled on all disk modules in the array. Each disk has internal write cache memory which is used to store data until it is written to the disk media. (This should not be confused with the array controller write cache, which is protected by the BBU.) If power to a disk is interrupted before data in write cache has been copied to the disk media, the data will be lost. To avoid this situation, the WCE feature should be disabled on all disks in the array. This ensures that data is written directly to the disk media and not stored in disk cache.

This procedure for disabling disk WCE should be executed any time a disk module is replaced or a new disk module is added to the array. This will ensure that WCE is disabled on all disks in the array.

To disable WCE on all disks, type:

**amutil -w <ArrayID>**

WCE will be disabled on all disks visible to the array controllers. Depending on the number of disk modules in the array, this command can take up to 60 seconds to complete.

### Command Example

The following example disables WCE on all disks in disk array 0000005EBD20:

**amutil -w 0000005EBD20**

## Enabling Disk Write Cache Enable (WCE)

***CAUTION*** *WCE should only be enabled in environments that provide uninterruptible power to the disk array. A loss of power to the disk array may result in data loss with WCE enabled.*

If maximum I/O performance is critical, disk WCE can be enabled on all the disks in the array. Disk WCE enhances disk array I/O performance, but increases the possibility of data loss during a power loss.

If WCE is enabled, this procedure for enabling disk WCE should be executed any time a disk module is replaced or a new disk module is added to the array. This will ensure that WCE is enabled on all disks in the array.

To enable WCE on all disks, type:

**amutil -w on <ArrayID>**

WCE will be enabled on all disks visible to the array controllers. Depending on the number of disk modules in the array, this command can take up to 60 seconds to complete.

### Command Example

The following example enables WCE on all disks in disk array ARRAY1:

**amutil -w on ARRAY1**

## Performing Disk Array Maintenance

At some point during operation of the disk array, you may need to perform maintenance tasks that involve error recovery and problem isolation. This section describes the tasks involved in maintaining the disk array.

### Locating Disk Modules

Array Manager 60 provides the means of identifying a disk module by flashing its amber Fault LED. You can flash the Fault LED on an individual disk, or on all the disks in the array. It is advisable to positively identify a failed disk before removing it from the disk array. Removing the wrong disk could cause the data on the LUN to become unavailable.

To locate a disk module, type:

**amutil <option> <ArrayID>**

- **option** identifies the disk modules you want to identify. Table 39 lists valid options.

**Table 39** Command Options for Locating Disk Array Components

| Option | Description |
|---|---|
| -f channel:ID | Flash the Fault LED on the specified disk. A disk is identified by both channel (1- 6) and SCSI ID (0 - 4, 8 - 12). |
| -F | Flash the Fault LEDs on all the disks in the disk array. |
| -s | Stop flashing the disk Fault LEDs. This option works for both a single disk and multiple disks. |

#### Command Example

The following example flashes the Fault LED on disk 3 on channel 2 in disk array data_center_1:

**amutil -f 2:3 data_center_1**

## Managing the Rebuild Process

If a disk fails, the disk array automatically begins the rebuild process the first time an I/O is performed to the LUN, providing that there is a global hot spare available. If no global hot spare is available, the rebuild will not occur until the failed disk has been replaced.

While a rebuild is in process, you can check its progress and change the rate at which the rebuild occurs. A rebuild must be in process to perform either of these tasks.

### Checking Rebuild Progress

To check the progress of a rebuild, type:

**amdsp -r <ArrayID>**

The progress of all rebuilds currently taking place on the disk array will be displayed.

### Changing Rebuild Priority Settings

The rebuild process must compete with host I/Os for disk array resources. Two settings allow you to control whether the rebuild process or host I/Os will have greater access to these resources. These settings reflect how important the rebuild process is relative to I/O activity on the LUN being rebuilt.

If you select settings that give higher priority to the rebuild process, it will complete sooner but at the cost of lower I/O performance. Lower priority settings give precedence to host I/Os, which may delay the completion of the rebuild. These settings are used only if host I/Os are occurring during the rebuild.

To change the rebuild priority settings, type:

**amutil -R <LUN> -f <freq> -a <amount> <ArrayID>**

- **freq** identifies the rate at which the disk array attempts to execute rebuild commands. Specified in tenths of a second, this value can be set to 1 to 50, or 0.1 seconds to 5.0 seconds. A low setting increases the frequency at which rebuild commands are issued, giving higher priority to the rebuild, but reducing I/O performance. A high value reduces the rebuild command frequency, giving higher priority to host I/Os. The default value for this setting is 1, or 0.1 seconds.

- **amount** identifies the number of blocks to rebuild at a time. This value can be from 1 to 64K and specifies the number of 512-byte blocks processed during each rebuild command. The higher the setting, the more blocks that will be processed, reducing I/O performance. A lower setting gives priority to host I/Os, delaying the completion of the rebuild. The default value for this setting is 64 blocks, or 32 Kbytes of data.

#### Command Example

The following example assigns a value of 5 seconds to the rebuild command rate, and sets the data block amount to 16 blocks on LUN 4 on disk array 0000005EBD20. This gives host I/Os higher priority than the default settings.

**amutil -R 4 -f 5 -a 16 0000005EBD20**

### Performing a Parity Scan

To verify the integrity of the parity data on a LUN, you can perform a parity scan. This will check each block of data against its parity to ensure that they match.

Like a rebuild, a parity scan competes with host I/Os for disk resources, and can impact host I/O performance. The rebuild priority **amount** setting is used when performing a parity scan. A larger amount value may have a greater impact on I/O performance. To avoid impacting performance, perform a parity scan during periods of low host activity.

To perform a parity scan, type:

**ammgr -P <LUN> <ArrayID>**

#### Command Example

The following example performs a parity scan on LUN 3 on disk array rack_1.

**ammgr -P 3 rack_1**

#### More About Parity Scanning

Parity information is used in RAID 5 LUNs to maintain data redundancy. When a single disk fails in a RAID 5 LUN, the disk array can reconstruct the data on the missing disk by using the parity information.

A parity scan compares the data and its associated parity to ensure they match. A parity scan cannot be performed on a LUN that has experienced a disk failure and is operating in degraded mode.

Although RAID 1 LUNs and 0/1 LUNs do not use parity, you can still perform a parity scan on them. The parity scan compares the data on the mirrored disks.

**Note** **What should I do if parity errors are detected?**
If errors are detected during a parity scan, it is recommended that you contact your Hewlett-Packard service representative immediately. The occurrence of parity errors may indicate a potential problem with the disk array hardware.

### Displaying Parity Scan Status

If a parity scan is in progress, you can monitor its progress. This will help you determine how much longer the parity scan will take.

To display parity scan progress, type:

**ammgr -s <LUN> <ArrayID>**

### Halting a Parity Scan

You can halt a parity scan currently in progress. This may be useful if the parity scan is impacting performance and you want to reschedule it to run when the host is less busy.

To halt a parity scan, type:

**ammgr -H <LUN> <ArrayID>**

## Managing Disk Array Logs

The disk array controller continuously gathers and stores information on all aspects of disk array operation. Array Manager 60 retrieves these controller log entries at regular intervals and stores them in a log directory on the host (`/var/opt/hparray/log`). The log files are valuable in monitoring disk array operation and isolating problems that may occur.

On HP08 and later, major event logging is available. MAjor event logging provides more detailed information on disk array operation. Many events that were not logged in the

previous firmware releases are logged in the major event log. Earlier versions of firmware (prior to HP08) use the standard log file format, also called Asychrnonous Event Notification (AEN).

**Note** On firmware HP08 and later, major event logging is enabled by default. If major event logging has been disabled by disabling the UTM, only standard log entries will be available.

### Allocating Space for Disk Array Logs

The amount of space available for storing log files on the host in controlled by the host environment variable `AM60_MAX_LOG_SIZE_MB`. This variable defines the maximum size (in megabytes) allocated for storage of the disk array log files. The log files are stored as individual files in /var/opt/hparray/log. All log files will be maintained until the specified maximum value is exceeded. At that point, the oldest log file will be deleted to make room for the new file (FIFO).

The valid range for `AM60_MAX_LOG_SIZE_MB` is 100 to 4096 (100 MB to 4 GB). Any value outside of this range is ignored and the default value is used. The default value for allocated log storage area is 100 MB.

**Note** - When changing the value for `AM60_MAX_LOG_SIZE_MB`, AM60Srvr must be stopped and restarted before the new value will take effect. To enforce the `AM60_MAX_LOG_SIZE_MB` value, it is necessary to stop and restart AM60Srvr every time the host is rebooted.

- When setting a large value for `AM60_MAX_LOG_SIZE_MB`, make sure the /var file system is large enough to accommodate the logs files. If it is not, the log files may overflow the `/var` file system.

### Viewing Disk Array Logs

To display the disk array controller log files, type:

**amlog [-s <StartTime>] [-e <EndTime>] [-t <Recordtype[,Recordtype].....>] [-c] [-d <LogDir>] [-a <ArrayID>]**

- **StartTime** identifies the starting date and time. Log entries with an earlier date and time will not be displayed. The default is the time of the oldest log entry. The format for entering the date and time is MMddhhmm[yy], where:

  MM = Month (01-12)
  dd = Day (01-31)
  hh = Hour (00-23)
  mm = Minute (00-59)
  yy = Year (00-99) [optional] Any number less than 90 will be evaluated as 2000 + yy

- **EndTime** identifies the ending date and time. Log entries with a later date and time will not be printed. The default is the time of the last log entry. Uses the same format as StartTime.

- **Recordtype** specifies the log record types to be displayed. The following record types are valid:

  – ctrlr - controller log sense. This is the default record type if none is specified.
  – mel - major event log. These entries can be filtered using the -c option. Major event logging is available in firmware versions HP08 and later.

- **-c** limits the major event log entries displayed to only those with critical status. Note that using this option will limit the amount of information displayed, including non-critical log entries that may be useful for isolating problems.

- **LogDir** identifies the path name of the log directory. The default is /var/opt/hparray/log The log entries are organized into a series of files with the naming convention AML000nn. When a log files reaches a certain size, a new file is created using the next sequential number. An associated catalog file (AMLCATLG) provides the information for navigating the binary entries in the main log files.

- **ArrayID** identifies a specific disk array for which log entries should be displayed. The default is to print entries for all disk arrays connected to and recognized by the host. The

actual ArrayID must be used here. An alias cannot be used because alias names are not recorded in the log file.

### Command Example

The following example displays the major event log entries for disk array rack_1. The log entries displayed are limited to only critical entries, and entries made after 0900 on 15 May 2000.

**amlog -s 05150000 -t mel -c rack_1**

## Sample Log Entries

The following is a sample of a major event log entry.

```
Major event log for Subsystem 001100A0B8060166 at Wed Nov 29
17:38:27 2000
   Sequence Number       = 14
   Time Stamp            = Wed Nov 29 17:38:19 2000
   Event Type            = Global Hot Spare Not Large Enough (224F)
   Event Category        = Notification
   Component Type        = Disk
   Enclosure             = 5
   Slot (0-based)        = 3
   Event Priority        = Informational
```

The following is a sample of an Asynchronous Event Log entry:

```
Controller log sense for Subsystem 001100A0B8060166 at Mon Dec  4
12:51:34 2000
   Controller Time Stamp = 032307 083318
   FRU Code              = 0x08
   FRU Code Qualifier    = 0x8142
   Sense Key             = 0x06
   Additional Sense Code = 0x3F
   Add Sense Code Qual   = 0xC7
Decoded Field Replaceable Unit Information:
   FRU Group             = Sub-Enclosure Group
   FRU ID Setting        = 1
   FRU Type              = Power Supply
```

```
   FRU State             = Failed
Decoded SCSI Sense:
Non-media Component Failure
   Reporting LUN         = 0
```

For information on interpreting SCSI sense codes, see "SCSI Sense Codes" on page 327.

### Flushing Disk Array Log Contents

Array Manager 60 automatically retrieves the contents of the disk array controller log at regular intervals, typically 15 minutes. However, if necessary you can manually flush (retrieve) the contents of the disk array log to the host. This may be useful if you suspect a very recent event has not yet been retrieved.

To flush the contents of the disk array log, type:

**amutil -l <ArrayID>**

A message is displayed indicating that the log file has been flushed. The log contents can now be displayed using the amlog command as described above.

### Purging Controller Logs

The controller logs are retrieved at regular intervals and stored on the host. Over time, these logs may grow quite large. Major event logging in particular generates very large log files. To reduce the space consumed by the logs, the oldest log file can be purged from the log file directory.

To purge the oldest log file in the host directory, type:

**amutil -p**

**Note** Always use the `amutil -p` command to purge the controller logs. This command maintains the catalog pointers used to access the log files. Using a system command such as `rm` to remove the log files will cause log catalog errors.

Management of the log files can be automated by creating a script that purges the oldest log files at regular intervals using `amutil -p`. This technique can be used to ensure that the log files don't grow to a size that may cause the ?var file system to overflow.

## Resetting Battery Age

The battery age should be reset to zero when the battery is replaced. This ensures an accurate indication of the age of the battery

To reset the battery age to zero, type:

**ammgr -b <ArrayID>**

## Installing Updated Patches

The following HP-UX patches are required when upgrading to HP08 firmware. The patches must be installed before upgraidng the contoller firmware to HP08. If the patches are not installed, the upgrade will fail.

- **HP-UX 10.20**: PHCO_22627 and PHSS_22846
- **HP-UX 11.0**: PHCO_22628 and PHSS_22847

The required patches can be downloaded from the following web sites:

http://us-support2.external.hp.com/index.html/ (External web site)
ftp://hpatlse.atl.hp.com/hp-ux_patches/ (Internal web site)

**Note** The patches are not currently included on the *HP-UX Support Plus CD-ROM*. They must be downloaded from the indicated web sites.

## Upgrading Disk Firmware

The firmware on each disk can be upgraded individually. Because different disks require different firmeware files, it may be necessary to

# Managing the Disk Array Using STM

STM is an online diagnostic tool, but it can be used to perform some of the common tasks involved in managing the disk array. The tasks described here are available to all users and do not require the purchase of a license. See "Support Tools Manager" on page 347 for more information on using this tool.

## Checking Disk Array Status Information

The STM Information Tool displays disk array status information. See "Using the STM Information Tool" on page 352 for more information on running and using this tool.

| STM Tool | Action |
|---|---|
| xstm, mstm | Select: Tools > Information Tool > Run |

## Binding a LUN

The STM Expert Tool can be used to bind a LUN. See "Using the STM Expert Tool" on page 355 for more information on running and using this tool.

| STM Tool | Action |
|---|---|
| xstm, mstm | Select: Tools > Expert Tool > Run<br>Select: Utilities > Bind LUN |

## Unbinding a LUN

The STM Expert Tool can be used to unbind a LUN. See "Using the STM Expert Tool" on page 355 for more information on running and using this tool.

| STM Tool | Action |
|---|---|
| xstm, mstm | Select: Tools > Expert Tool > Run<br>Select: Utilities > Unbind LUN |

## Adding a Global Hot Spare

The STM Expert Tool can be used to add a global hot spare. See "Using the STM Expert Tool" on page 355 for more information on running and using this tool.

| STM Tool | Action |
|---|---|
| xstm, mstm | Select: Tools > Expert Tool > Run<br>Select: Utilities > Hot Spares > Create |

## Removing a Global Hot Spare

The STM Expert Tool can be used to remove a global hot spare. See "Using the STM Expert Tool" on page 355 for more information on running and using this tool.

| STM Tool | Action |
|---|---|
| xstm, mstm | Select: Tools > Expert Tool > Run<br>Select: Utilities > Hot Spares > Delete |

## Locating Disk Modules

The STM Expert Tool can be used to locate disk modules. to aid in identification. The LEDs on the disk array components are flashed to aid in identification. See "Using the STM Expert Tool" on page 355 for more information on running and using this tool.

| STM Tool | Action |
|---|---|
| xstm, mstm | Select: Tools > Expert Tool > Run<br>Select: Utilities > Flash Component Lights > Drive<br>or<br>Utilities > Flash Component Lights > LUN<br>or<br>Utilities > Flash Component Lights > Disk Enclosure<br>or<br>Utilities > Flash Component Lights > Array |

# Status Conditions and Sense Code Information

The following tables may be useful interpreting the various types of disk array status information that is returned by the management tools. Where appropriate, any required action is identified.

## LUN Status Conditions

The LUN status condition terminology used by Array Manager 60 (AM60) may differ from that used by STM. Both terms are identified in the table.

**Table 40** LUN Status Conditions

| Status | Definition/Action |
|---|---|
| **AM60**: OPTIMAL<br>**STM**: OPTIMAL | The LUN is operating normally.<br>No action is required. |
| **AM60**: OPTIMAL--PARITY SCAN IN PROGRESS<br>**STM**: OPTIMAL | A parity scan is in progress on the disk array.<br>No action is required. |
| **AM60**: OPTIMAL-REDUNDANCY RE-INITIALIZATION PENDING OR IN PROGRESS<br>**STM**: OPTIMAL | A background parity sync operation is in progress on this LUN.<br>No action is required. The LUN can be used immediately. |
| **AM60**: DEGRADED--WAITING FOR REPAIR ACTION<br>**STM**: DEGRADED - 1 | A disk has failed and there is no hot spare to perform a rebuild.<br>Replace the failed disk or add a hot spare. |
| **AM60**: DEGRADED--REPLACED DISK BEING FORMATTED<br>**STM**: DEGRADED - 65 | A disk has been replaced and is being formatted prior to starting a rebuild on this LUN.<br>No action is required. |

**Table 40** LUN Status Conditions (cont'd)

| Status | Definition/Action |
|---|---|
| **AM60**: DEGRADED--REPLACED DISK BEING REBUILT<br>**STM**: DEGRADED - 2 | A rebuild is in progress on the LUN.<br>No action is required. |
| **AM60**: DEAD--MORE DISK FAILURES THAN REDUNDANT DISKS<br>**STM**: UNAVAILABLE - 4 | Multiple, simultaneous disk failures have occurred on the LUN, causing data to be inaccessible. On a RAID 5 LUN, losing more than one disk will cause this status. On a RAID 1 or 0/1 LUN, you must lose both disks of a mirrored pair before this status will occur.<br>Replace the failed disks. |
| **AM60**: DEAD--FORMAT IN PROGRESS<br>**STM**: BINDING | A formatting operation (writing zeros) is in progress on a DEAD LUN. This occurs while the LUN is being bound.<br>No action is required. |
| **AM60**: UNDEFINED<br>**STM**: DEGRADED - 83 OR UNAVAILABLE - 83 | This state indicates a drive failure during a LUN binding process. If a hot spare has been designated, the LUN state should transition to Optimal once the LUN binding is complete. If no hot spare is in place, the LUN state should transition to Degraded--Waiting for Repair Action (Degraded--1), once the binding is complete. |
| **AM60**: DEAD--CREATION IN PROGRESS<br>**STM**: UNAVAILABLE - 84 | A formatting operation has been queued for this LUN, but has not started yet.<br>No action is required. |
| **AM60**: DEAD--WRONG DISK REMOVED OR REPLACED<br>**STM**: UNAVAILABLE - 116 | The wrong disk was removed from a degraded LUN.<br>Reinstall the disk. |
| **AM60**: INACCESSIBLE:POSSIBLY NO DEVICE FILE<br>**STM**: UNAVAILABLE - 254 | The special device files are not installed on the LUN.<br>Run insf -e on the LUN. |
| **AM60**: UNDEFINED<br>**STM**: UNAVAILABLE - xxx | The LUN is in an unknown state.<br>Contact your service representative. |

## Disk Status Conditions

The disk status condition terminology used by Array Manager 60 (AM60) may differ from that used by STM. Both terms are identified in the table.

**Table 41** Disk Status Conditions

| Status | Definition/Action |
|---|---|
| **AM60:** OPTIMAL<br>**STM:** OPTIMAL (OPT) | The disk is operating normally.<br>No action is required. |
| **AM60:** NON-EXISTENT<br>**STM:** No Disk (NIN) | The disk array controller has no knowledge of a disk in this slot. This may be caused by no disk being installed, or a disk that never responds to a SCSI command from the controller.<br>No action is required. |
| **AM60:** UNASSIGNED<br>**STM:** OPTIMAL (OPT) | The disk is a member of the unassigned disk group. It is available for use in a LUN or as a hot spare.<br>No action is required. |
| **AM60:** FAILED--CAUSE UNKNOWN<br>**STM:** FLT - 3 | The disk has failed due to an unknown condition.<br>Replace the disk. |
| **AM60:** REPLACED<br>**STM:** OFF - 4 | The disk has been replaced. The disk will maintain this status until the original data has been rebuilt or copied to the disk.<br>No action is required. |
| **AM60:** WRONG DRIVE REMOVED OR REPLACED<br>**STM:** FLT - 9 | The wrong disk was removed or replaced from a degraded LUN. The new disk gets this status.<br>Reinstall the original disk. |
| **AM60:** OFFLINE--INTERNALLY MARKED OUT OF SERVICE<br>**STM:** OFF - 11 | A disk has been installed that was part of another LUN, possibly in another disk array. The disk array attempts to recreate the original LUN, but because some of the disk are missing, the disk is marked OFFLINE. The newly created LUN will also be assigned a status of DEAD.<br>Unbind the LUN the disk array has created. This will return the disk to UNASSIGNED status. |
| **AM60:** NON-SUPPORTED CHANNEL<br>**STM:** FLT - 17 | The command issued specified an invalid channel.<br>Retry the command using the proper channel. |

**Table 41** Disk Status Conditions (cont'd)

| Status | Definition/Action |
|---|---|
| **AM60:** READ FAILED<br>**STM:** FLT - 19 | The disk array could not read from the disk.<br>Replace the failed disk. |
| **AM60:** WRONG BLOCK SIZE<br>**STM:** OFF - 22 | The disk uses an incompatible block size (not 512 bytes).<br>Replace with a supported disk. |
| **AM60:** DISK LOCKED OUT<br>**STM:** UNSUPP | The disk is not supported.<br>Install a supported disk. |
| **AM60:** NON-SUPPORTED ID<br>**STM:** FLT - 33 | The command made a request using an unsupported ID.<br>Retry the command with the proper parameters. |
| **AM60:** NO RESPONSE OR DEFINED GLOBAL HOT SPARE--DRIVE IS MISSING<br>**STM:** FLT - 35 | The disk is not responding. It may have failed or been removed.<br>Replace or reinstall the disk. |
| **AM60:** CAPACITY LESS THAN MINIMUM<br>**STM:** OFF - 38 | A disk was replaced by a disk with lower capacity.<br>Install disk with the proper capacity.<br><br>The disk is unsupported.<br>Install a supported disk. |
| **AM60:** NON-SUPPORTED CHANNEL & ID<br>**STM:** FLT - 49 | The command made a request to an unsupported channel and ID.<br>Retry the command using the proper channel & ID. |
| **AM60:** FORMAT OR RECONSTRUCT FAILED<br>**STM:** FLT - 51 | The disk failed while being formatted or rebuilt.<br>Replace the disk. |
| **AM60:** WRITE FAILED<br>**STM:** FLT - 67 | The disk array could not write to the disk.<br>Replace the failed disk. |
| **AM60:** UNDEFINED<br>**STM:** FLT - xxx | The disk is an unknown state. |
| **AM60:** FAILED VIA MODE SELECT<br>**STM:** FLT - 83 | The disk was failed using a host-initiated command. |

## Component Status Conditions

Component status conditions are organized into the categories listed in Table 42. The interpretation and action associated with a status depends on the component. See Table 51 on page 379 for more information on Disk System SC10 component status.

The component status condition terminology used by Array Manager 60 (AM60) may differ from that used by STM. Both terms are identified in the table.

Table 42 Component Status Conditions

| Status | Description |
|---|---|
| **AM60:** Good<br>**STM:** Optimal | The component is operating normally. |
| **AM60:** Non-critical<br>**STM:** Fault - 3 | A component or condition such as temperature or voltage has exceeded warning limits.<br><br>Although immediate action is not required, the cause of the non-critical status should be investigated and corrected if necessary. |
| **AM60:** Critical<br>**STM:** Fault - 2 | A component has failed or a condition such as temperature or voltage has exceeded critical limits. Immediate action is required. |
| **AM60:** Not Installed<br>**STM:** Not Present | The component is not installed. |
| **AM60:** Unknown<br>**STM:** Fault - 0 | The status of the component cannot be determined. |

## FRU Codes

The FRU codes indicate which disk array component is responsible for the log entry. Log entries that do not involve disk modules typically require you to interpret the FRU Code and the FRU Code Qualifier values to determine which component is identified.

To simplify reporting events, components within the disk array have been placed in FRU groups. The FRU Code indicates which FRU group the component is a member of. The FRU Code Qualifier further identifies the specific component within the group, and the component's status.

**Table 43** FRU Code Groups

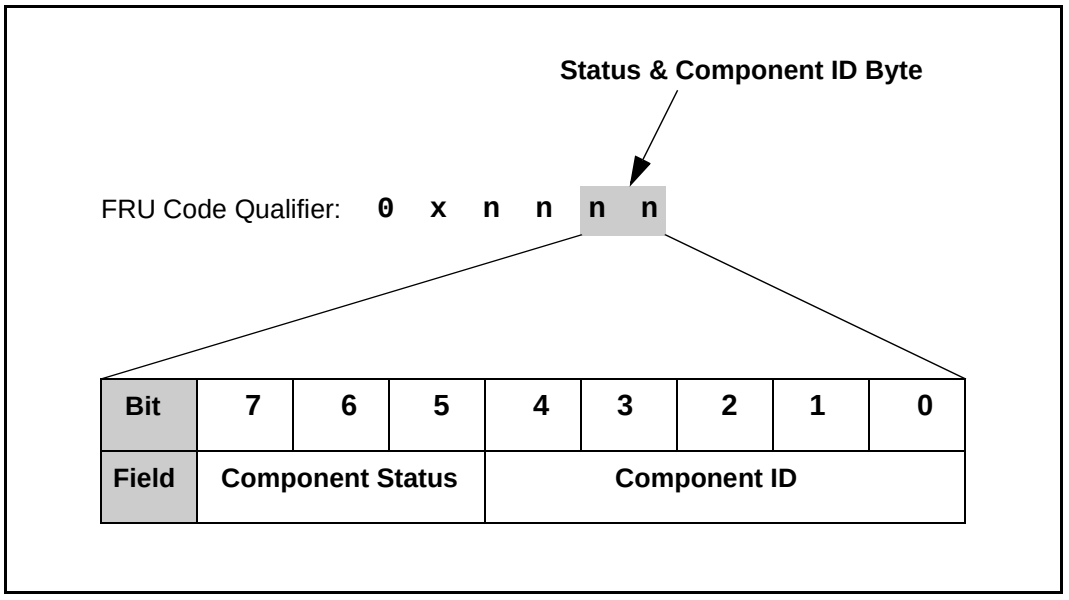
| FRU Code Value | Group Description |
|---|---|
| 0x01 | Host Channel Group - comprises the host SCSI bus, its SCSI interface chip, and all initiators and other targets connected to the bus. |
| 0x02 | Controller Drive Interface Group - comprises the SCSI interface chips on the controller which connect to the drive buses. |
| 0x03 | Controller Buffer Group - comprises the controller logic used to implement the on-board data buffer. |
| 0x04 | Controller Array ASIC Group - comprises the ASICs on the controller associated with the array functions. |
| 0x05 | Controller Other Group - comprises all controller related hardware not associated with another group. |
| 0x06 | Controller Enclosure Group - comprises components that are monitored by the array controller, such as power supplies, fans, thermal sensors, and AC power monitors. See "Controller Enclosure Group FRU Code Qualifier" on page 324 for information on identifying component and status. |
| 0x07 | Subsystem Configuration Group - comprises subsystem components that are configurable by the user, on which the array controller will display information (such as faults). |

**Table 43** FRU Code Groups

| FRU Code Value | Group Description |
|---|---|
| 0x08 | Disk Enclosure Group - comprises attached disk enclosures. This group includes the power supplies, environmental monitor, and other components in the disk enclosure. See "Disk Enclosure Group FRU Code Qualifier" on page 326 for information on identifying component and status. |
| 0x09-0x0F | Reserved |
| 0x10-0xFF | Drive Groups - comprises a drive (embedded controller, drive electronics, and Head Disk Assembly), its power supply, and the SCSI cable that connects it to the controller; or supporting sub-enclosure environmental electronics. The FRU code designates the channel ID in the most significant nibble and the SCSI ID of the drive in the least significant nibble.<br><br>NOTE: Channel ID 0 is not used because a failure of drive ID 0 on this channel would cause an FRU code of 0x00, which the SCSI-2 standard defines as no specific unit has been identified to have failed or that the data is not available. |

## Controller Enclosure Group FRU Code Qualifier

When the Controller Enclosure group is identified (FRU Code = 0x06), the FRU Code Qualifier is interpreted as follows:

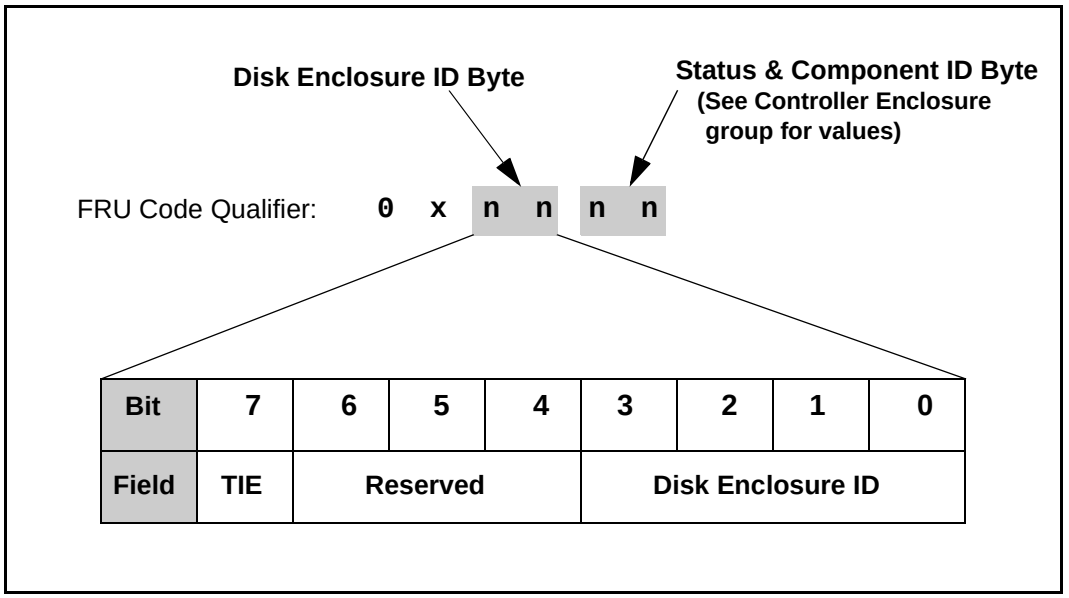



| Component Status Value | Status |
|---|---|
| 0 | Optimal |
| 1 | Warning |
| 2 | Failed |
| 3 | Missing |

| Component ID Value | Component |
|---|---|
| 0 | Unspecified |
| 1 | Device |
| 2 | Power Supply |
| 3 | Cooling Element |
| 4 | Temperature Sensors |
| 6 | Audible Alarm |
| 7 | Environmental Services Electronics |
| 8 | Controller Electronics |
| 9 | Nonvolatile Cache |
| B | Uninterruptible Power Supply |
| 0x0C - 0x13 | Reserved |
| 0x14 | SCSI Target Port |
| 0x15 | SCSI Initiator Port |

## Disk Enclosure Group FRU Code Qualifier

When the Disk Enclosure group is identified (FRU Code = 0x08), the FRU Code Qualifier is interpreted as follows:




| TIE Value | Disk Enclosure ID |
|---|---|
| 0 | When TIE (Tray Identifier Enable) is set to 0, the Disk Enclosure ID field indicates both the channel and enclosure as follows:<br><br>`          Enclosure #1     Enclosure #2`<br>`Channel 1     0x01             0x06`<br>`Channel 2     0x02             0x07`<br>`Channel 3     0x03             0x08`<br>`Channel 4     0x04             0x09`<br>`Channel 5     0x05             0x0A` |
| 1 | When TIE is set to 1, the Disk Enclosure ID field value corresponds to the enclosure (tray) ID set on the disk enclosure. |

## SCSI Sense Codes

Table 44 lists the SCSI sense codes that may be returned as part of the log entry. This information may be helpful interpreting log entries. Only the Additional Sense Code and Additional Sense Code Qualifier fields are required to identify each condition.

**Table 44** SCSI Sense Codes

| Additional Sense Code | Additional Sense Code Qualifier | Interpretation |
|---|---|---|
| 00 | 00 | No Additional Sense Information<br><br>The controller has no sense data available for the requesting host and addressed logical unit combination. |
| 04 | 01 | Logical Unit is in the Process of Becoming Ready<br><br>The controller is executing its initialization functions on the addressed logical unit. This includes drive spinup and validation of the drive/ logical unit configuration information. |
| 04 | 02 | Logical Unit Not Ready, Initializing Command Required<br><br>The controller is configured to wait for a Start Stop Unit command before spinning up the drives, but the command has not yet been received. |
| 04 | 04 | Logical Unit Not Ready, Format In Progress<br><br>The controller previously received a Format Unit command from an initiator, and is in the process of executing that command. |
| 04 | 81 | Storage Module Firmware Incompatible - Manual Code Synchronization Required |
| 04 | A1 | Quiescence Is In Progress or Has Been Achieved |

**Table 44** SCSI Sense Codes (cont'd)

| Additional Sense Code | Additional Sense Code Qualifier | Interpretation |
|---|---|---|
| 0C | 00 | **If the accompanying sense key = 4, error is interrupted as follows:**<br><br>Unrecovered Write Error<br>Data could not be written to media due to an unrecoverable RAM, battery or drive error.<br><br>**If the accompanying sense key = 6, error is interrupted as follows:**<br><br>Caching Disabled<br>Data caching has been disabled due to loss of mirroring capability or low battery capacity. |
| 0C | 01 | Write Error Recovered with Auto Reallocation<br><br>The controller recovered a write operation to a drive and no further action is required by the host. Auto reallocation may not have been used, but this is the only standard ASC/ ASCQ that tells the initiator that no further actions are required by the driver. |
| 0C | 80 | Unrecovered Write Error Due to Non-Volatile Cache Failure<br><br>The subsystem Non-Volatile cache memory recovery mechanisms failed after a power cycle or reset. This is possibly due to some combination of battery failure, alternate controller failure or foreign controller.<br><br>User data may have been lost. |
| 0C | 81 | Deferred Unrecoverable Error Due to Memory Failure<br><br>Recovery from a Data Cache error was unsuccessful.<br><br>User data may have been lost. |

**Table 44** SCSI Sense Codes (cont'd)

| Additional Sense Code | Additional Sense Code Qualifier | Interpretation |
|---|---|---|
| 11 | 00 | Unrecovered Read Error<br><br>An unrecovered read operation to a drive occurred and the controller has no redundancy to recover the error (RAID 0, degraded RAID 1, degraded mode RAID 3, or degraded RAID 5). |
| 11 | 8A | Miscorrected Data Error - Due to Failed Drive Read<br><br>A media error has occurred on a read operation during a reconfiguration operation,<br><br>User data for the LBA indicated has been lost. |
| 18 | 02 | Recovered Data - Data Auto Reallocated<br><br>The controller recovered a read operation to a drive and no further action is required by the host. Auto reallocation may not have been used, but this is the only standard ASC/ ASCQ that tells the initiator that no further actions are required by the driver. |
| 1A | 00 | Parameter List Length Error<br><br>A command was received by the controller that contained a parameter list and the list length in the CDB was less than the length necessary to transfer the data for the command. |
| 20 | 00 | Invalid Command Operation Code<br><br>The controller received a command from the initiator that it does not support. |

**Table 44** SCSI Sense Codes (cont'd)

| Additional Sense Code | Additional Sense Code Qualifier | Interpretation |
|---|---|---|
| 21 | 00 | Logical Block Address Out of Range<br><br>The controller received a command that requested an operation at a logical block address beyond the capacity of the logical unit. This error could be in response to a request with an illegal starting address or a request that started at a valid logical block address and the number of blocks requested extended beyond the logical unit capacity. |
| 24 | 00 | Invalid Field in CDB<br><br>The controller received a command from the initiator with an unsupported value in one of the fields in the command block. |
| 25 | 00 | Logical Unit Not Supported<br><br>The addressed logical unit is currently unconfigured. An Add LUN operation in the Logical Array Mode Page must be executed to define the logical unit before it is accessible. |
| 26 | 00 | Invalid Field in Parameter List<br><br>The controller received a command with a parameter list that contained an error. Typical errors that return this code are unsupported mode pages, attempts to change an unchangeable mode parameter, or attempts to set a changeable mode parameter to an unsupported value. |
| 28 | 00 | Not Ready to Ready Transition<br><br>The controller has completed its initialization operations on the logical unit and it is now ready for access. |
| 29 | 00 | Power On, Reset, or Bus Device Reset Occurred<br><br>The controller has detected one of the above conditions. |

**Table 44** SCSI Sense Codes (cont'd)

| Additional Sense Code | Additional Sense Code Qualifier | Interpretation |
|---|---|---|
| 29 | 04 | Device Internal Reset<br><br>The controller has reset itself due to an internal error condition. |
| 29 | 81 | Default Configuration has been Created<br><br>The controller has completed the process of creating a default logical unit. There is now an accessible logical unit that did not exist previously. The host should execute its device scan to find the new logical unit. |
| 29 | 82 | Controller Firmware Changed Through Auto Code Synchronization<br><br>The controller firmware has been changed through the Auto Code Synchronization (ACS) process. |
| 2A | 01 | Mode Parameters Changed<br><br>The controller received a request from another initiator to change the mode parameters for the addressed logical unit. This error notifies the current initiator that the change occurred. This error may also be reported in the event that Mode Select parameters changed as a result of a cache synchronization error during the processing of the most recent Mode Select request. |
| 2A | 02 | Log Parameters Changed<br><br>The controller received a request from another initiator to change the log parameters for the addressed logical unit. This error notifies the current initiator that the change occurred. This error is returned when a Log Select command is issued to clear the AEN log entries. |

**Table 44** SCSI Sense Codes (cont'd)

| Additional Sense Code | Additional Sense Code Qualifier | Interpretation |
|---|---|---|
| 2F | 00 | Commands Cleared by Another Initiator<br><br>The controller received a Clear Queue message from another initiator. This error is to notify the current initiator that the controller cleared the current initiators commands if it had any outstanding. |
| 31 | 01 | Format Command Failed<br><br>A Format Unit command issued to a drive returned an unrecoverable error. |
| 32 | 00 | Out of Alternates<br><br>A Re-assign Blocks command to a d rive failed. |
| 3F | 01 | Drive micro-code changed |
| 3F | 0E | Reported LUN's data has changed.<br>Previously LUN data reported via a Report LUNs command has changed (due to LUN creation/deletion or controller hot swap) |

**Table 44** SCSI Sense Codes (cont'd)

| Additional Sense Code | Additional Sense Code Qualifier | Interpretation |
|---|---|---|
| 3F | 8N | Drive No Longer Usable.<br><br>The controller has set a drive to a state that prohibits use of the drive. The value of N in the ASCQ indicates the reason why the drive cannot be used.<br><br>0 - The controller set the drive state to "Failed - Write failure"<br>1 - Not used<br>2 - The controller set the drive state to "Failed" because it was unable to make the drive usable after replacement. A format or reconstruction error occurred.<br>3 - Not used<br>4 - Not used<br>5 - The controller set the drive state to "Failed - No response"<br>6 - The controller set the drive state to "Failed - Format failure"<br>7 - The controller set the drive state to "User failed via Mode Select"<br>8 - Not used<br>9 - The controller set the drive state to "Wrong drive removed/ replaced"<br>A - Not used<br>B - The controller set the drive state to "Drive capacity < minimum"<br>C - The controller set the drive state to "Drive has wrong block size"<br>D - The controller set the drive state to "Failed - Controller storage failure"<br>E - Drive failed due to reconstruction failure at Start of Day (SOD) |
| 3F | 98 | Drive Marked Offline Due to Internal Recovery Procedure<br><br>An error has occurred during interrupted write processing causing the LUN to transition to the DEAD state. Drives in the drive group that did not experience the read error will transition to the Offline state (0x0B) and log this error. |

**Table 44** SCSI Sense Codes (cont'd)

| Additional Sense Code | Additional Sense Code Qualifier | Interpretation |
|---|---|---|
| 3F | BD | The controller has detected a drive with Mode Select parameters that are not recommended or which could not be changed. Currently this indicates the QErr bit is set incorrectly on the drive specified in the FRU field of the Request Sense data. |
| 3F | C3 | The controller had detected a failed drive side channel specified in the FRU Qualifier field. |
| 3F | C7 | Non-media Component Failure<br><br>The controller has detected the failure of a subsystem component other than a disk or controller. The FRU codes and qualifiers indicate the faulty component. |
| 3F | C8 | AC Power Fail<br><br>The Uninterruptible Power Source has indicated that AC power is no longer present and the UPS has switched to standby power. |
| 3F | C9 | Standby Power Depletion Imminent<br><br>The UPS has indicated that its standby power source is nearing depletion. The host should take actions to stop IO activity to the controller. |
| 3F | CA | Standby Power Source Not at Full Capability<br><br>The UPS has indicated that its standby power source is not at full capacity. |
| 3F | CB | AC Power Has Been Restored<br><br>The UPS has indicated that AC power is now being used to supply power to the controller. |

**Table 44** SCSI Sense Codes (cont'd)

| Additional Sense Code | Additional Sense Code Qualifier | Interpretation |
|---|---|---|
| 3F | D0 | Write Back Cache Battery Has Been Discharged<br><br>The controller's battery management has indicated that the cache battery has been discharged. |
| 3F | D1 | Write Back Cache Battery Charge Has Completed<br><br>The controller's battery management has indicated that the cache battery is operational. |
| 3F | D8 | Cache Battery Life Expiration<br><br>The cache battery has reached the specified expiration age. |
| 3F | D9 | Cache Battery Life Expiration Warning<br><br>The cache battery is within the specified number of weeks of failing. |
| 3F | E0 | Logical Unit Failure<br><br>The controller has placed the logical unit in a "dead" state. User data and/ or parity can no longer be maintained to ensure availability. The most likely cause is the failure of a single drive in non-redundant configurations or a second drive in a configuration protected by one drive. The data on the logical unit is no longer accessible. |
| 3F | EB | LUN marked DEAD due to Media Error Failure during SOD<br><br>An error has occurred during interrupted write processing causing the LUN to transition to the DEAD state |

**Table 44** SCSI Sense Codes (cont'd)

| Additional Sense Code | Additional Sense Code Qualifier | Interpretation |
|---|---|---|
| 40 | NN | Diagnostic Failure on Component NN (0x80 - 0xFF)<br><br>The controller has detected the failure of an internal controller component. This failure may have been detected during operation as well as during an on-board diagnostic routine. The values of NN supported in this release of the software are listed below.<br><br>80 - Processor RAM<br>81 - RAID Buffer<br>82 - NVSRAM<br>83 - RAID Parity Assist (RPA) chip or cache holdup battery<br>84 - Battery Backed NVSRAM or Clock Failure<br>91 - Diagnostic Self Test failed non-data transfer components test<br>92 - Diagnostic SelfTest failed data transfer components test<br>93 - Diagnostic SelfTest failed drive Read/ Write Buffer data turnaround test<br>94 - Diagnostic Self Test failed drive Inquiry access test<br>95 - Diagnostic Self Test failed drive Read/ Write data turnaround test<br>96 - Diagnostic Self Test failed drive Self Test |
| 43 | 00 | Message Error<br><br>The controller attempted to send a message to the host, but the host responded with a Reject message. |

**Table 44** SCSI Sense Codes (cont'd)

| Additional Sense Code | Additional Sense Code Qualifier | Interpretation |
|---|---|---|
| 44 | 00 | Internal Target Failure<br><br>The controller has detected a hardware or software condition that does not allow the requested command to be completed.<br><br>**If the accompanying sense key is 0x04:**<br>Indicates a hardware failure. The controller has detected what it believes is a fatal hardware or software failure and it is unlikely that a retry would be successful.<br><br>**If the accompanying sense key is 0x0B:**<br>Indicates an aborted command. The controller has detected what it believes is a temporary software failure that is likely to be recovered if retried. |
| 45 | 00 | Selection Time-out on a Destination Bus<br><br>A drive did not respond to selection within a selection time-out period. |
| 47 | 00 | SCSI Parity Error<br><br>The controller detected a parity error on the host SCSI bus or one of the drive SCSI buses. |
| 48 | 00 | Initiator Detected Error Message Received<br><br>The controller received an Initiator Detected Error Message from the host during the operation. |
| 49 | 00 | Invalid Message Error<br><br>The controller received a message from the host that is not supported or was out of context when received. |

**Table 44** SCSI Sense Codes (cont'd)

| Additional Sense Code | Additional Sense Code Qualifier | Interpretation |
|---|---|---|
| 49 | 80 | Drive Reported Reservation Conflict<br><br>A drive returned a status of reservation conflict. |
| 4B | 00 | Data Phase Error<br><br>The controller encountered an error while transferring data to/ from the initiator or to/ from one of the drives. |
| 4E | 00 | Overlapped Commands Attempted<br><br>The controller received a tagged command while it had an untagged command pending from the same initiator or it received an untagged command while it had a tagged command(s) pending from the same initiator. |
| 5D | 80 | Disk SMART Event (Self-Monitoring Analysis and Reporting Technology) |
| 80 | 02 | Bad ASC code detected by Error/ Event Logger |
| 80 | 03 | Error occurred during data transfer from SRM host. |
| 84 | 00 | Operation Not Allowed With the Logical Unit in its Current State<br><br>The requested command or Mode Select operation is not allowed with the logical unit in the state indicated in byte 76 of the sense data. Examples would be an attempt to read or write a "dead" logical unit or an attempt to verify or repair parity on a "degraded" logical unit. |
| 84 | 06 | LUN Awaiting Format<br><br>A mode select has been done to create a LUN but the LUN has not been formatted. |

**Table 44** SCSI Sense Codes (cont'd)

| Additional Sense Code | Additional Sense Code Qualifier | Interpretation |
|---|---|---|
| 85 | 01 | Drive IO Request Aborted<br><br>IO Issued to Failed or Missing drive due to recently failed removed drive. This error can occur as a result of I/Os in progress at the time of a failed or removed drive. |
| 87 | 00 | Microcode Download Error<br><br>The controller detected an error while downloading microcode and storing it in non-volatile memory. |
| 87 | 08 | Incompatible Board Type For The Code Downloaded |
| 87 | 0C | Download failed due to UTM LUN number conflict |
| 87 | 0E | Controller Configuration Definition Inconsistent with Alternate Controller |
| 88 | 0A | Subsystem Monitor NVSRAM values configured incorrectly. |
| 8A | 00 | Illegal Command for Drive Access<br><br>The initiator attempted to pass a command through to a drive that is not allowed. The command could have been sent in pass-thru mode or by attempting to download drive microcode. |
| 8A | 01 | Illegal Command for the Current RAID Level<br><br>The controller received a command that cannot be executed on the logical unit due to its RAID level configuration. Examples are parity verify or repair operations on a RAID 0 logical unit. |
| 8A | 10 | Illegal Request- Controller Unable to Perform Reconfiguration as Requested<br><br>The user requested a legal reconfiguration but the controller is unable to execute the request due to resource limitations. |

**Table 44** SCSI Sense Codes (cont'd)

| Additional Sense Code | Additional Sense Code Qualifier | Interpretation |
|---|---|---|
| 8B | 02 | Quiescence Is In Progress or Has Been Achieved |
| 8B | 03 | Quiescence Could Not Be Achieved Within the Quiescence Timeout Period |
| 8B | 04 | Quiescence Is Not Allowed |
| 8E | 01 | A Parity/ Data Mismatch was Detected<br><br>The controller detected inconsistent parity/data during a parity verification. |
| 91 | 00 | General Mode Select Error<br><br>An error was encountered while processing a Mode Select command. |
| 91 | 03 | Illegal Operation for Current Drive State<br><br>A drive operation was requested through a Mode Select that cannot be executed due to the state of the drive. An example would be a Delete Drive when the drive is part of a LUN. |
| 91 | 09 | Illegal Operation with Multiple SubLUNs Defined<br><br>An operation was requested that cannot be executed when multiple SubLUNs are defined on the drive. |
| 91 | 33 | Illegal Operation for Controller State<br><br>The requested Mode Select operation could not be completed due to the current state of the controller. |

**Table 44** SCSI Sense Codes (cont'd)

| Additional Sense Code | Additional Sense Code Qualifier | Interpretation |
|---|---|---|
| 91 | 36 | Command Lock Violation<br><br>The controller received a Write Buffer Download Microcode, Send Diagnostic, or Mode Select command, but only one such command is allowed at a time and there was another such command active. |
| 91 | 3B | Improper Volume Definition for Auto-Volume Transfer mode - AVT is disabled.<br><br>Controller will operate in normal redundant controller mode without performing Auto-Volume transfers. |
| 91 | 50 | Illegal Operation For Drive Group State<br><br>An operation was requested that cannot be executed due to the current state of the Drive Group. |
| 91 | 51 | Illegal Reconfiguration Request - Legacy Constraint<br><br>Command could not be completed due to Legacy configuration or definition constraints |
| 91 | 53 | Illegal Reconfiguration Request - System Resource Constraint<br><br>Command could not be completed due to resource limitations of the controller |
| 94 | 01 | Invalid Request Due to Current Logical Unit Ownership |
| 95 | 01 | Extended Drive Insertion/Removal Signal<br><br>The controller has detected the drive insertion/removal signal permanently active. |

**Table 44** SCSI Sense Codes (cont'd)

| Additional Sense Code | Additional Sense Code Qualifier | Interpretation |
|---|---|---|
| 95 | 02 | Controller Removal/Replacement Detected or Alternate Controller Released from Reset<br><br>The controller detected the activation of the signal/signals used to indicate that the alternate controller has been removed or replaced. |
| 98 | 01 | The controller has determined that there are multiple sub-enclosures with the same ID value selected. |
| 98 | 02 | Sub-enclosure with redundant ESMs specifying different Tray levels. |
| 98 | 03 | Sub-enclosure ESMs have different firmware levels |
| A0 | 00 | Write Back Caching Could Not Be Enabled<br><br>The controller could not perform write back caching due to a battery failure or discharge, Two Minute Warning signal from the UPS, or an ICON failure. |
| A1 | 00 | Write Back Caching Could Not Be Enabled - RDAC Cache Size Mismatch<br><br>The controller could not perform write back caching due to the cache sizes of the two controllers in the RDAC pair not matching. |
| A4 | 00 | Global Hot Spare Size Insufficient for All Drives in Subsystem.<br><br>A defined Global Hot Spare is not large enough to cover all of the drives present in the subsystem. Failure of a drive larger than the Global Hot Spare will not be covered by the Global Hot Spare drive. |

**Table 44** SCSI Sense Codes (cont'd)

| Additional Sense Code | Additional Sense Code Qualifier | Interpretation |
|---|---|---|
| A6 | 00 | Recovered processor memory failure<br><br>The controller has detected and corrected a recoverable error in processor memory. |
| A7 | 00 | Recovered data buffer memory error<br><br>The controller has detected and corrected a recoverable error in the data buffer memory.<br><br>Sense bytes 34-36 will contain the count of errors encountered and recovered |
| C0 | 00 | The Inter-controller Communications Have Failed<br><br>The controller has detected the failure of the communications link between redundant controllers. |
| D0 | 06 | Drive IO Time-out<br><br>The controller destination IO timer expired while waiting for a drive command to complete. |
| D1 | 0A | Drive Reported Busy Status<br><br>A drive returned a busy status in response to a command. |
| E0 | XX | Destination Channel Error<br><br>XX = 00 through 07 indicates the Sense Key returned by the drive after a check condition status<br><br>XX = 10 indicates that a bus level error occurred |

**Table 44** SCSI Sense Codes (cont'd)

| Additional Sense Code | Additional Sense Code Qualifier | Interpretation |
|---|---|---|
| E0 | 20/21 | Fibre Channel Destination Channel Error<br><br>ASCQ = 20: Indicates redundant path is not available to devices<br>ASCQ = 21: Indicates destination drive channels are connected to each other<br><br>Sense byte 26 will contain the Tray ID<br>Sense byte 27 will contain the Channel ID |

# 5 HP-UX DIAGNOSTIC TOOLS

Overview . . . . . 346

Support Tools Manager . . . . . 347

# Overview

STM (Support Tools Manager) is the primary diagnostic tool available for the Disk Array FC60. For diagnosing problems, STM provides the capability to gather and display detailed status information about the disk array. STM can also be used to perform common management tasks.

# Support Tools Manager

The STM host-based utility provides capability for managing the Disk Array FC60. STM comes with HP-UX instant ignition and support media.

The Support Tools Manager (STM) host-based utility is the primary online diagnostic tool available for the HP SureStore E Disk Array FC60. STM provides the capability for testing, configuring, and evaluating the operational condition of the disk array. STM comes with HP-UX instant ignition and support media.

**Note** Using the information in this chapter requires some understanding of the basic operation of STM. If you are not familiar with STM, go to the Systems Hardware, Diagnostics, and Monitoring section of HP's Online Documentation web site (www.docs.hp.com), which contains information on using STM, including a tutorial to get you started. You can also download a copy of the Diagnostic/IPR Media User's Guide, which contains complete information on using STM.

## STM User Interfaces

STM presents three different user three interfaces: an X-Windows interface (xstm), a menu-based interface (mstm), and a command-based interface (cstm). Each STM interface has an extensive online help facility as well as man pages.

**Note** Each Disk Array FC60 is displayed twice in the STM interface — once for each active disk array controller. Each controller is identified by a unique hardware path. STM operations can be performed by selecting either disk array controller. The effect on the disk array will be the same regardless of which controller is selected.

STM identifies the Disk Array FC60 by its product number of A5277A. This is the identifier you will look for when selecting a disk array in STM.

## xstm — the X Windows Interface

xstm is the X-Windows screen-based STM interface. Because it is the easiest to use, xstm is the recommended interface for systems that support graphical displays.

The main xstm window displays a map representing system resources. The STM system map represents each Disk Array FC60 as two icons labeled “A5277A Array”. See figure Figure 89. Each icon represents one of the disk array controllers, which are identified by their hardware paths. Click on the icon for the disk array you will be testing.

Navigation through xstm is done using menus and submenus. For more information on navigating in xstm, download a copy of the STM Quick Reference Guide from the Systems Hardware, Diagnostics, and Monitoring section of HP’s Online Documentation website (www.docs.hp.com).

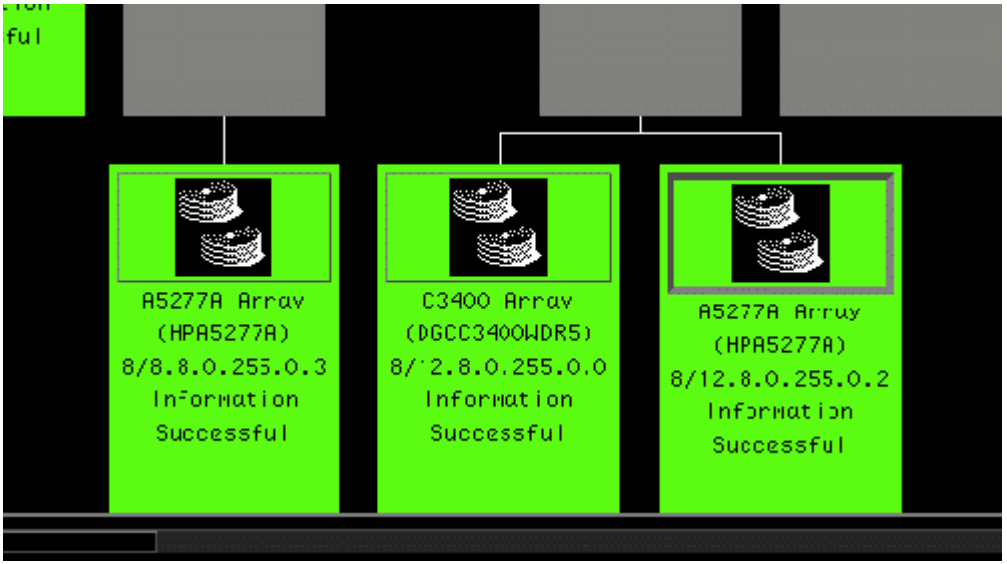



**Figure 89** xstm Interface Main Window Disk Array Icons

### mstm — the Menu-based Interface

mstm is the menu-based STM interface. It serves as an alternate interface for systems that do not support graphical displays.

The main mstm window displays a list of system resources. The Disk Array FC60 is identified as product type “A5277A Array”. See Figure 90. Each entry in the list represents one of the disk array controllers, which are identified by their hardware paths. Select the entry for the disk array you will be testing.

In mstm, you traverse screens and menus, and select commands from pulldown menus, which are similar to those found in xstm.

- **Navigating between screens.** To navigate between one screen and another, use the Tab key.
- **Navigating within screens.** To navigate from one portion to another portion of the same screen , use the Prev and Next keys (or, alternatively, the Cursor Up and Cursor Down keys).
- **Navigating the menu bar.** The following methods are used to navigate the menu bar:
  - Position the cursor on the menu bar — Use the Tab key (or the MenuBar on/off function key).
  - Move to a particular pulldown menu — Use the cursor arrow keys.
  - Expand a menu/sub-menu — Use the Return key.
  - Highlight a command or sub-menu — Use the cursor keys.
  - Perform a command — Use the Return key.
  - Invoke a menu directly — Use Alt function key; then press nonunderlined character in menu title.
  - Activate a menu command/expand pulldown sub-menu — Press the letter that is underlined in command or submenu.

For more information on navigating in mstm, download a copy of the STM Quick Reference Guide from the Systems Hardware, Diagnostics, and Monitoring section of HP’s Online Documentation website (www.docs.hp.com/hpux/systems).

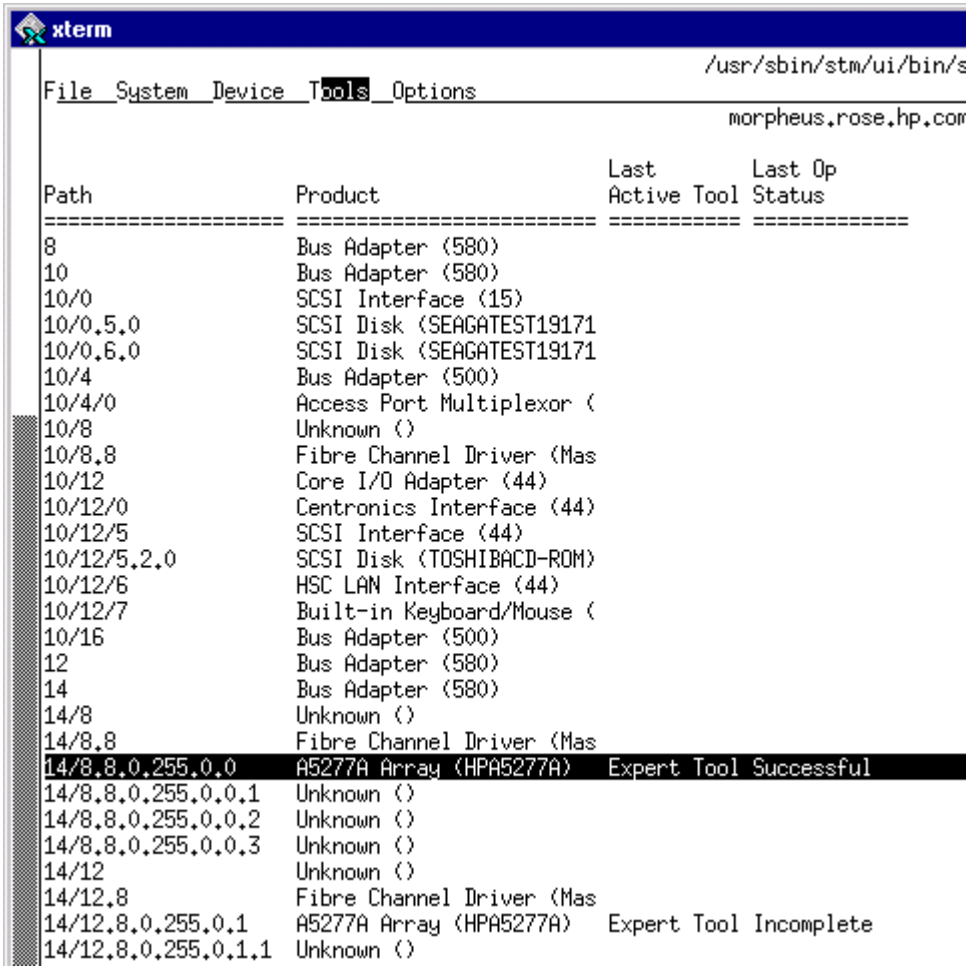



**Figure 90** mstm Interface Main Window

## STM Tools

The STM tools available for use with the HP SureStore E Disk Array FC60 are listed in Table 45.

**Table 45** Available Support Tools

| Tool type | Description |
|---|---|
| Information | Provides detailed configuration and status information for all disk arrays components. |
| Expert | Provides capability to perform common disk array management tasks. |

### STM Device Logs

When you run an STM support tool, the results are indicated by the status of the device icon on the system map (xstm) or the device entry on the device list (mstm, cstm). If the results are other than successful, you can view the appropriate device log for more information.

- Failure Log - If the results indicate Failure, view the device Failure Log for information identifying the likely cause of the failure and recommended action for correcting the problem.
- Test Activity Log - If the results are other than Successful or Failure, view the Test Activity Log for an explanation of what occurred.

## Using the STM Information Tool

The STM Information Tool gathers status and configuration information about the selected disk array and stores this information in three logs: the information log, the activity log, and the failure log.

### Running Information Tool in X Windows

1. At the system prompt:
   - Type **xstm &**
2. Click on the desired A5277A Array device icon.
3. To run the Information Tool and view the Information log:
   - From the menu bar, select Tools
   - Select Information
   - Select Run. The Information Tool builds and displays the Information log.
   - Select Done when done viewing the log.
4. To view the Activity log:
   - From the menu bar, select Tools
   - Select Information
   - Select Activity Log.
   - Select Done when done viewing the log.
5. To exit STM:
   - From the menu bar, select File
   - Select Exit

### Running Information Tool in Menu Mode

1. At the system prompt:
   - Type **mstm**
   - Select Ok
2. To select the desired disk array**:**
   - Scroll down using the arrow key, select the A5277A Array
   - Press <Enter>.
3. To run the Information Tool and display the Information log:
   - From the Menubar, select Tools
   - Select Information
   - Select Run. The Information Tool builds and displays the Information log.
   - Select Done when done viewing the log
4. To view the Activity log:
   - From the Menubar, select Tools
   - Select Information
   - Select Activity Log
   - Select Done when done viewing the log
5. To exit STM:
   - From the Menubar, select File
   - Select Exit
   - Select OK

## Interpreting the Information Tool Information Log

The Information Log contains status and configuration information for all disk array components. The log is separated into the following sections:

- Controller Enclosure – information for the components in the disk array controller enclosure
- Disk Enclosure(s) – information about the components in the disk enclosure(s), excluding the disks
- Map – a graphic layout showing the configuration of the disks enclosures and disks
- Configuration – information on LUN configuration
- Disks – information about each disk

Much of the information is self-explanatory, but if you need assistance interpreting any fields, a complete online description is available by performing the following steps:

1. Run the Information tool on the desired disk array as described in the preceding section.
2. Return to the main STM window and select the Tools menu.
3. Select Information.
4. Select Info. The Information Log Description window is displayed describing each field.
5. Select Done to close the description window.

## Interpreting the Information Tool Activity Log

The Information Tool Activity Log allows you to review details of the most recent execution of the Information tool against a specific device. You can view, save, and print the contents of the activity log, as desired.

## Using the STM Expert Tool

The Expert Tool provides the capability to manage the HP SureStore E Disk Array FC60. Before using the Expert Tool for the first time you are encouraged to read through the Expert Tool help topics. The Step-by-Step instructions in particular provide useful tips on using the Expert Tool.

As you perform tasks using the Expert Tool, the status of each operation is displayed in the main window. The Expert Tool window also provides a status of the activities performed on the selected device.

### Running Expert Tool in X Windows

1. At the system prompt:
   - Type **xstm &**
2. Click on the A5277A ARRAY device icon
3. To run the Expert tool:
   - From the menu bar, select Tools
   - Select Expert
   - Select Run
4. Perform the desired operation using the menus in the Expert Tool window. The Expert Tool menu options are listed in Table 46.
5. To exit the Expert Tool:
   - From the menu bar, Select File
   - Select Exit

### Running Expert Tool in Menu Mode

1. At the system prompt:
   - Type **mstm**
   - Select Ok
2. To select the disk array:
   - Scroll down using the arrow key, select A5277A ARRAY.
   - Press <Enter>
3. To run the Expert tool:
   - Select Menubar on or use arrow keys to get to Menubar
   - Select Tools
   - Select Expert Tool
   - Select Run
4. Perform the desired operation using the Expert Tool menus. The Expert Tool menu options are listed in Table 46.
5. To exit Expert Tool:
   - Select Menubar on or use arrow keys to get to Menubar
   - Select File
   - Select Exit
   - Select Yes

### Expert Tool Menu Options (xstm and mstm)

When using xstm or mstm, all Expert Tool functions are accessed from a series of menus. The menus options are described in Table 46.

**Table 46** Expert Tool Menus and Descriptions

| Menu | Option | Property | Description |
|---|---|---|---|
| Logs | View Event Log | NA | Displays selected event log entries |
| Tests | Parity Scan | NA | Perform a parity scan on a LUN. |
| Utilities | Bind LUN | NA | Bind selected disk modules into a LUN with a specified RAID level. |
| | Unbind LUNs | NA | Unbind a LUN. |
| | Replace LUN Zero | NA | Unbind and rebind LUN 0 |
| | Hot Spares | Create | Create hot spares. |
| | | Delete | Delete hot spares. |
| | Flash Component Lights | Drive | Flash Fault LED on selected disk. |
| | | LUN | Flash Fault LEDs on all disks in the selected LUN. |
| | | Disk Enclosure | Flash Fault LEDs on all disks in the selected disk enclosure. |
| | | Array | Flash Fault LEDs on all disks in the disk array. |

# 6 TROUBLESHOOTING

Introduction . . . . . 360

Disk Array Installation/Troubleshooting Checklist . . . . . 365

Power-Up Troubleshooting . . . . . 366

Controller Enclosure Troubleshooting. . . . . 367

SureStore E Disk System SC10 Troubleshooting . . . . . 376

# Introduction

The modular design of the Disk Array FC60 simplifies the isolation and replacement of failed hardware components. Most disk array components are hot-swappable Field Replaceable Units (FRUs), which can be replaced while the disk array is operating. Some of the FRUs are customer replaceable. Other array components can be replaced in the field, but only by a trained service representative. A complete list of product and part numbers are included in "Replaceable Parts" on page 418.

***CAUTION*** *To ensure proper operation, disk array controller modules and disk modules should not be removed or replaced when disk array power is off. Also, disk modules should not be moved to different slot locations within the disk enclosure.*

## About Field Replaceable Units (FRUs)

The Disk Array FC60 consists of a Controller Enclosure and one or more SureStore E Disk System SC10 enclosures. Table 47 identifies the disk array FRUs and whether they are customer replaceable. See "Removal and Replacement" on page 383 for more information.

**Table 47** Field Replaceable Units

| Field Replaceable Units | Customer Replaceable Unit | Hot Swappable Modules |
|---|---|---|
| **Disk Enclosure** | | |
| Disk Modules | Yes | Yes |
| BCC (Bus Controller Card) Modules<br>– Note - disk enclosure should be powered down when replacing the BCC. | No | |
| Power Supply Modules | Yes | Yes |
| Disk Fan Modules | Yes | Yes |
| **Controller Enclosure** | | |
| Controller Modules | No | Yes |
| Media Interphase Adapter (MIA) | Yes | Yes |
| Battery Backup Module | Yes | Yes |
| Controller Fan Module | Yes | Yes |
| SCSI Cables | Yes | No |
| Fibre Channel Cables | Yes | Yes |
| Power Supply Modules | Yes | Yes |
| Power Supply Fan Module | Yes | Yes |

## HP-UX Troubleshooting Tools

There are several tools available for troubleshooting the disk array on an HP-UX host. This includes monitoring the operation of the disk array and gathering information that will help identify and solve the problem.

- Array Manager 60 - primarily used to manage the disk array, Array Manager 60 can also be used to check the status of the disk array and to retrieve log information.
- Support Tools Manager (STM) - the Information Tool of STM is one the most valuable sources of information for identifying problems with the disk array.

## Windows NT and Windows 2000 Troubleshooting Tools

The *HP Storage Manager 60* software includes a set of tools for troubleshooting the disk array. See the *HP Storage Manager 60-NT User's Guide* for information on using the *HP Storage Manager 60* software.

## EMS Hardware Event Monitoring (HP-UX Only)

EMS hardware monitoring is one the most important troubleshooting tools available for the Disk Array FC60. With EMS monitoring, you can be alerted to problems as they occur, allowing you to respond quickly to correct a problem before it impacts disk array operation. All components and aspects of disk array operation are monitored. And EMS gives you the flexibility to deliver event notification using a variety of different methods.

More information on EMS Hardware Monitoring can be downloaded from the Systems Hardware, Diagnostics, and Monitoring section of HP's Online Documentation web site at www.docs.hp.com/hpux/systems.

## EMS Monitor Event Severity Levels

Each event detected and reported by the EMS monitor is assigned a severity level, which indicates the impact the event may have on disk array operation. The following severity levels are used for all events:

| | |
|---|---|
| Critical | An event that causes host system downtime, or other loss of service. Host system operation will be affected if the disk system continues to be used without correction of the problem. Immediate action is required. |
| Serious | An event that may cause, host system downtime, or other loss of service if left uncorrected. Host system and hardware operation may be adversely affected. The problem needs repair as soon as possible. |
| Major Warning | An event that could escalate to a serious condition if not corrected. Host system operation should not be affected and normal use of the disk system can continue. Repair is needed but at a convenient time. |
| Minor Warning | An event that will not likely escalate to a severe condition if left uncorrected. Host system operation will not be interrupted and normal use of the disk system can continue. The problem can be repaired at a convenient time. |
| Information | An event that is expected as part of the normal operation of the hardware. No action is required. |

## Sample EMS Event Message

A sample EMS hardware monitor event message is shown in Figure 91. A message typically includes the following information:

- Message Data – Date and time the message was sent, the source and destination of the message, and the severity level
- Event Data – Date and time of the event, the host, event ID, name of the monitor, event number, event class, severity level, hardware path, associated OS error log entry ID
- Error Description – Information indicating the component that experienced the event and the nature of the event

- Probable Cause/Recommended Action – The cause of the event and suggested steps toward a solution. This information should be the first step in troubleshooting.

```
Notification Time:  Thu Aug 6 15:18:03 1998

yourserver sent Event Monitor notification information:

/peripherals/events/mass_storage/LVD/enclosure/10_12.8.0.255.0.10.0 is
!=1.
Its current value is CRITICAL(5)

Event data from monitor:

Event Time: Thu Aug 6 15:18:03 1998
Hostname: yourserver.rose.hp.com      IP Address: 15.43.212.175
Event ID: 0x0035ca2b9b00000002       Monitor: dm_ses_enclosure
Event # : 101                                        Event Class: I/O
Severity: Critical

Enclosure at hardware path 10/12.8.0.255.0.10.0: Hardware failure
Associated OS error log entry id(s): None

Description of Error:

     The disk device is indicating a fault.

Probable Cause/Recommended Action:

    The disk drive in slot 1 has failed. Check the connection, reseat,
or replace the disk drive.
```

**Figure 91** Sample EMS Hardware Event Notification Message

# Disk Array Installation/Troubleshooting Checklist

The following checklist is intended to help isolate and solve problems that may occur when installing the disk array.

- Check Fibre Optic and SCSI Cables and SCSI Terminators:
  - No damaged fibre optic cables
  - No damaged or loose screws on connectors
  - All cables tightly secured to the connectors on the Fibre Channel Interconnect PCA
  - Shortest possible fibre optic cable lengths between disk arrays and host adaptors
- Check Disk Array (Pre-Power-Up):
  - All modules properly seated in the disk array enclosures
  - Disk modules in their proper slots
- Check Disk Array Functionality and Configuration (Following power up)
  - Date and time synchronized with the host.
  - No fault indicators on any disk array components
  - All disk modules spun up
  - LUN 0 replaced on non-factory built disk array (A5277A).
  - Host ioscan displays proper disk array configuration
  - Fibre Channel drivers installed
  - Different capacity disk modules not mixed in same LUN
  - Disk modules with different part/product numbers not mixed in same LUN

# Power-Up Troubleshooting

When the disk array is powered up, each component perform an internal self-test, to ensure it is operating properly.

Visual indications of power-up are:

- The green Power LED on the controller enclosure is on
- The green Power LED on each disk enclosures is on
- All fans are operating
- No Fault LEDs are on. See Figure 92 on page 368 and Figure 93 on page 377.

The following LEDs will also be on:

- Controller enclosure:
  - The green Full Charge LEDs on the controller batteries, if the batteries are fully charged. If the batteries are not fully charged, the Full Charge LEDs will flash until charging is complete, and then remain on.
  - The Power LEDs on the controller enclosure
- Disk enclosure
  - The green Power Supply LEDs
  - Fan LEDs

If the power-up sequence fails, the appropriate amber Fault LED will be on, identifying the failing component. If the power-up tests fail, all modules in the disk array must be examined for Fault LEDs that are on. See Table 48 on page 369 and Table 50 on page 377).

**Note** Inserting a module while the power-up tests are running can cause the tests to fail. Insert all modules and then cycle power to rerun the power-up tests.

**Note** If no LEDs are ON and the fans are not running, it indicates that no AC power is being supplied to the disk array power supply modules. Check the input AC power to the disk array.

See "Applying Power to the Disk Array" on page 198 for information on powering up the disk array.

# Controller Enclosure Troubleshooting

## Introduction

This section contains information on identifying problems with the disk array controller enclosure.

## Controller Enclosure LEDs

Figure 92 shows the locations of the status LEDs for the controller enclosure. Table 48 summarizes the operating LED states for all components within the controller enclosure.

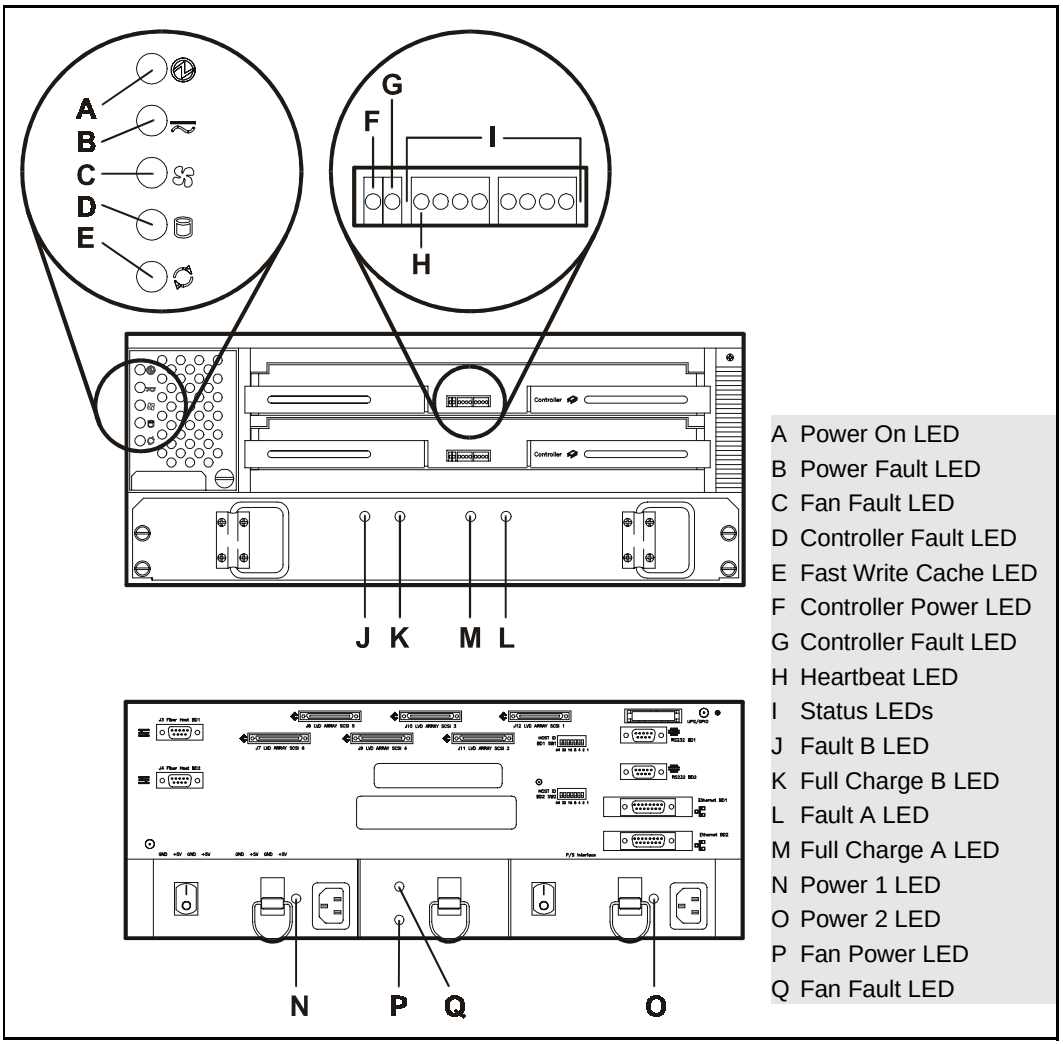



**Figure 92** Controller Enclosure LEDs

Table 48 Normal LED Status for Controller Enclosure

| Module | LED | Normal State |
|---|---|---|
| Controller Enclosure | Power On | On (green) |
| | Power Fault | Off |
| | Fan Fault | Off |
| | Controller Fault | Off |
| | Fast Write Cache | On (green) while data is in cache |
| Controller | Controller Power | On (green) |
| | Controller Fault | Off |
| | Heartbeat | Blink (green) |
| | Status | Green<br>There are 8 status LEDs. The number and pattern of these LEDs depend on how your system is configured. |
| Controller Battery | Fault - B | Off |
| | Full Charge - B | On (green)[1] |
| | Fault - A | Off |
| | Full Charge - A | On (green)[1] |
| Controller Power Assembly | Power 1 | On (green) |
| | Power 2 | On (green) |
| Controller Fan Assembly | Fan Power | On (green) |
| | Fan Fault | Off |

1. Both Full Charge A and Full Charge B LEDs are ON after batteries are fully charged. The LEDs flash while charging is in progress, and remain on when charging is complete.

## Master Troubleshooting Table

Table 49 contains troubleshooting information for the controller enclosure and modules.

Table 49 Master Troubleshooting Controller

| Symptom | Possible Cause | Procedure |
|---|---|---|
| **Controller** | | |
| Controller LED (front cover) is on and the Fan Fault LED is off. | A Controller missing or unplugged | Check the power LEDs on both controller modules. If a Power LED is off, make sure that the module is plugged in correctly and its handles are locked in place. |
| | B Controller failed | If the Fault LED remains on after replacing the controller, go to cause C. |
| | C One or more memory modules failed | Replace the memory modules.<br>If the Fault LED remains ON after replacing the memory, go to cause D. |
| | D Controller enclosure midplane failed | Replace the midplane. If the Fault LED remains ON after replacing the midplane, call the factory service center. |
| Software issued a controller error message | Controller failed | Check the Fan Fault LED on the front cover. If it is on, go to Troubleshooting Controller fan module Problems. If not, continue at the next step. Replace the failed controller. |
| Controller enclosure and fan fault LEDs (front cover) are on | Both a controller enclosure and the controller fan module failed. | Check both controllers for fault LEDs, then replace the failed controller.<br>After replacing the failed controller, check the LEDs on the front cover. If the Fan Fault LED is still on, replace the controller module. |

Table 49 Master Troubleshooting Controller (cont'd)

| Symptom | Possible Cause | Procedure |
|---|---|---|
| Controller enclosure and Fan Fault LED (front cover) are on | Controller enclosure fan failure caused one or both controller(s) to overheat | 1. Stop all activity to the controller module and turn off the power.<br>2. Replace the failed controller enclosure fan module.<br>3. Allow the controller to cool down, then turn on the power.<br>4. Check both controllers for fault LEDs. If a Controller Fault LED is on, replace the failed controller. |
| **Midplane** | | |
| Power LEDs (front and power supply FRUs) are ON but all other Power LEDs are off | Other FRUs are missing or not installed correctly. | Check all FRUs in the controller enclosure and make sure they are installed securely. If this does not solve the problem, go to cause B. |
| | Defective DC power harness or power interface board. | Replace the DC power harnesses and the power supply assembly. If this does not solve the problem, go to cause C. |
| | Failed controller enclosure midplane. | Replace the controller enclosure midplane. |

Table 49 Master Troubleshooting Controller (cont'd)

| Symptom | Possible Cause | Procedure |
|---|---|---|
| Software errors occur when attempting to access controller or disks | A Software function or configuration problems | Check the appropriate software and documentation to make sure the system is set up correctly or that the proper command was executed. |
| | B Controller enclosure power switches or main circuit breakers in rackmount cabinet turned off | Make sure that all power switches are turned on. |
| | C Loose, disconnected, or defective interface cables | Check all the cables between the host, controller enclosure, disk enclosure, and other devices on the network. Make sure they are undamaged (particularly fiber optic cables) and securely attached. |
| | D Improper ID settings | Make sure that the Fibre host IDs and SCSI disk enclosure IDs are set properly for all devices, particularly if you just removed or added a new device to the network. |
| | E Failed controller enclosure or disk | Check the controller enclosure for a controller module fault. If the Controller LED on the front cover is turned on, look at each controller module's LED for a fault.<br>Check the disk enclosure for a disk fault. If a failed disk is indicated, replace the disk. |
| | F Failed controller enclosure midplane | Replace the controller enclosure midplane. |

Table 49 Master Troubleshooting Controller (cont'd)

| Symptom | Possible Cause | Procedure |
|---|---|---|
| **Controller Fan Module** | | |
| Fan Fault LED is on | One or both of the fans in the controller fan module has failed. | Replace the controller fan module. |
| | The power supply fan module is unplugged or has failed. | 1. Make sure the power supply fan module is plugged in correctly. Reseat the module if necessary.<br>2. Check the LEDs on the power supply fan module. If the Power LED is off or the Fan Fault LED is on, replace the power supply fan module.<br><br>If the Power LED is on and the Fan Fault LED is off, the problem may be caused by the controller fan module. |
| | The temperature sensor in the controller fan module has detected an over temperature condition. | Even if the fans in the controller module are operating, an obstructed air flow path or a high ambient temperature can cause the temperature within the enclosure to rise. Determine what is causing the temperature problem and correct it. |
| No LEDs ON | Controller fan module is unplugged or has failed. | 1. Make sure the controller fan module is plugged in correctly. Reseat the module if necessary.<br>2. If the Power LED remains off, replace the controller power module. |
| **Battery** | | |
| Fault-A or Fault-B LED (BBU) is on | Left or right battery bank has failed. | Replace the BBU. |
| Full Charge-A or Full Charge-B LED (BBU) is off | Left or right battery bank is not yet fully charged. | Turn ON the power and allow controller module to run 7 hours to recharge the batteries.<br>If after 7 hours, the battery low error persists, replace the BBU. |

Table 49 Master Troubleshooting Controller (cont'd)

| Symptom | Possible Cause | Procedure |
|---|---|---|
| "Battery Low" error issued by software | Power turned OFF for extended period and drained battery power. | Turn ON the power and allow controller module to run 7 hours to recharge the batteries.<br>If after 7 hours, the battery low error persists, replace the BBU. |
| | Batteries are weak and FRU is due for replacement. | Check the last service date for the BBU. If it has been two years since the battery was replaced, replace the BBU. |
| BBU is new but will not hold a charge | Battery charger board failure | Replace the BBU. Allow the system to run for at least 7 hours in order to properly charge the batteries. If this does not solve the problem, go to cause B. |
| | Faulty battery harness | Replace both power supplies.<br>Replace the battery harness. |
| Loss of battery backup during a power outage | Battery failure or BBU is missing. | Replace the BBU. |
| **Power Supply Module** | | |
| No power to the controller module (all Power LEDs off) | Power switches are turned off. | Turn ON both power switches on the controller enclosure.<br>Turn ON the main circuit breakers in the rack, if applicable. |
| | Power cords unplugged. | Make sure all power cords are plugged in securely. |
| | Power supply modules overheated or failed. | Lower the room temperature if it is too high, unblock enclosure vents if they are obstructed, replace any failed fans in the enclosure, or replace the power supply module. |
| One Power Supply LED (FRU) is ON and the other is off | A Power supply module is turned OFF or unplugged. | Plug in the power cord and turn ON the power switch. If the LED remains off, go to cause B. |
| | B Power supply module is overheated or failed. | Lower the room temperature if it is too high, unblock enclosure vents if they are obstructed, replace any failed fans in the enclosure, or replace the power supply module. |

Table 49 Master Troubleshooting Controller (cont'd)

| Symptom | Possible Cause | Procedure |
|---|---|---|
| Power Supply LED (front cover) is on | A Power supply module is missing or not plugged in properly. | Insert and lock the power supply module into place. If the Fault LED is still on, go to cause B. |
| | B Power supply module is overheated or failed. | Lower the room temperature if it is too high, unblock enclosure vents if they are obstructed, replace any failed fans in the enclosure, or replace the power supply module. |
| The LEDs on both power supply modules are ON but all other LEDs are off. | A Defective DC Power Harness | – Replace the DC Power Harnesses. If this does not resolve the problem, go to cause B. |
| | B Defective Power Interface Board | – Replace the power supply. |

# SureStore E Disk System SC10 Troubleshooting

This section contains information on identifying and isolating problems with the Disk System SC10 disk enclosure.

## Disk Enclosure LEDs

Figure 93 shows the locations of the disk enclosure status LEDs. Table 50 summarizes the operating LED states for all components within the disk enclosure.

## Losing LUN 0

If LUN 0 becomes unavailable because of multiple disk failures, Array Manager 60 may not be able to communicate with the disk array. If this occurs, replace the failed disks as quickly as possible. When the disk are replaced it may be necessary to replace LUN 0, recreating it in its original configuration.

***CAUTION*** *If LUN 0 is lost, do not power off the disk array. Without LUN 0, the disk array will not operate normally when it is powered back on.*

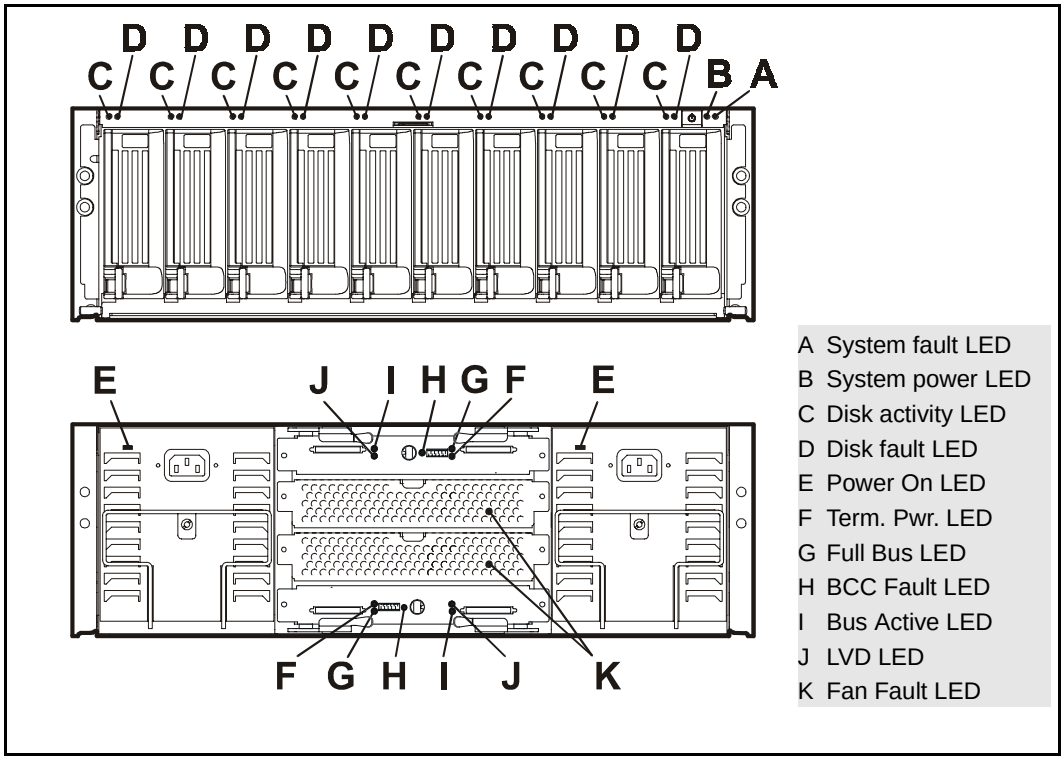

D D D D D D D D D D
C C C C C C C C C C B A
E J I H G F E
F G H I J K
A System fault LED
B System power LED
C Disk activity LED
D Disk fault LED
E Power On LED
F Term. Pwr. LED
G Full Bus LED
H BCC Fault LED
I Bus Active LED
J LVD LED
K Fan Fault LED

**Figure 93** Disk Enclosure LEDs

**Table 50** Disk Enclosure LED Functions

| LED | State | Indication |
|---|---|---|
| System Power | Green | Power is on. Normal operation. |
| | OFF | Power is off |
| System Fault | Amber | Self-test[1] / Problem[2] |
| | OFF | Normal operation |
| | Flashing | BCC A & B DIP switch settings do not match |

**Table 50** Disk Enclosure LED Functions (cont'd)

| LED | State | Indication |
|---|---|---|
| BCC Fault | Amber | Self-test[1] / Fault |
| | OFF | Normal operation |
| | Flashing | Peer BCC DIP switch settings do not match |
| LVD | Green | Bus operating in LVD mode |
| | OFF | Bus operating in single-ended mode |
| Term. Pwr. | Green | Termination power is available from the host. Normal operation. |
| | OFF | There is no termination power. |
| Full Bus | Green | Disks are on a single bus of ten addresses. |
| | OFF | Disks are split between two buses, five addresses each |
| Bus Active | Green | Bus is available for use. |
| | OFF | Isolator chip is disabled and bus is not available. |
| Fan | Green | Normal operation |
| | Amber | Startup[1] / Fault |
| | OFF | Power is off |
| Power Supply | Green | Normal operation |
| | Amber | Startup[1] / Fault |
| | OFF | Power is off |
| Disk Fault | OFF | Disk is operating normally. |
| | Amber | A fault has occurred with the disk. |
| | Flashing | Used to identify a disk using the software management tools. |
| Disk Activity[3] | Flashing | Indicates that the disk is being accessed by the disk array. |

1 Startup and self-tests occur briefly when the unit is powered ON.
2 A component has failed; temperature or voltage is out of normal range.
See "Isolating Causes" on page 380
3 When a disk is installed with power on, its activity LED stays ON until the disk has spun up. When the disk is ready, the LED turns off. Thereafter, it flashes when there is I/O to the disk.

**Note** It is normal for the amber Fault LED on a component to go on briefly when the component initially starts up. However, if the Fault LED remains on for more than a few seconds, a fault has been detected.

## Interpreting Component Status Values (HP-UX Only)

Common status terms have specific indications for various disk enclosure components. The component status condition terminology used by Array Manager 60 (AM60) may differ from that used by STM. Both terms are identified in Table 51.

**Table 51** Status Indications by Component

| Component/Status | Disks | BCCs | Fans, Power Supplies | Sensors |
|---|---|---|---|---|
| **AM60:** Good<br>**STM:** Optimal | Component is installed and no error conditions are known. | | | |
| **AM60:** Critical<br>**STM:** Fault - 2 | N/A | Hardware has failed. | | Voltage/temperature exceeds critical limit. |
| **AM60:** Non-critical<br>**STM:** Fault - 3 | N/A | BCC A and B have different firmware versions. | N/A | Voltage/temperature exceeds warning limit. |
| **AM60:** Not Installed<br>**STM:** Not Present | Component is not installed. | | | |
| **AM60:** Unknown<br>**STM:** Fault - 0 | N/A | N/A | N/A | Sensor has failed or status is not available. |
| Not Available | Component is installed without known errors, but has not been turned on or set into operation; or the controller that owns reporting has failed. | | | |

## Isolating Causes

Table 52 lists the probable causes and solutions for problems you may detect on the disk enclosure. When more than one problem applies to your situation, investigate the first description that applies. The table lists the most basic problems first and excludes them from subsequent problem descriptions.

**Table 52** Disk Enclosure Troubleshooting Table

| Problem Description | HW Event Category | LED State | Status | Probable Cause/Solution |
|---|---|---|---|---|
| The disk enclosure does not power on | none | System Power LED off | none | – Neither power cord is plugged in.<br>– The power switch is not pressed.<br>– AC breaker is tripped.<br>– AC power source has failed.<br>– The PDU/PDRU is defective.<br>– Power switch is defective.<br>– A faulty component is causing power supplies to turn off. Remove all components and reinsert one at a time until the faulty component is isolated. |
| System fault LED is on | none | Power Supply LED off | Power supply Not Available | – The power supply is not plugged in.<br>– The PDU/PDRU or primary power source has failed. |
| | Critical | Part Fault LED on | Critical | A component has failed. See problem descriptions below. |
| Buzzer sounds when BCC installed | Critical | BCC Fault flashing | Critical | BCC DIP switch settings do not match peer BCC switch settings. Reset switches. |
| BCC Fault LED is on | Critical | BCC Fault on | Critical | BCC hardware is faulty. Replace the BCC. |
| Fan LED is amber | Critical | Amber | Critical | Fan has slowed or stopped. Replace the fan. |

**Table 52** Disk Enclosure Troubleshooting Table (cont'd)

| Problem Description | HW Event Category | LED State | Status | Probable Cause/Solution |
|---|---|---|---|---|
| Power Supply LED is amber | Critical | Amber | Critical | – An incompatible or defective component caused a temporary fault.<br>– Power supply hardware is faulty.<br>Unplug the power cord and wait for the LED to turn off. Reinsert the power cord. If fault persists, replace the power supply. |
| Temperature is over limit | Critical | none | Critical Temp is >54.5º C | – A fan is faulty. Check status and correct.<br>– Airflow is obstructed; vents are blocked.<br>– One or more slots are empty.<br>– Power supply is faulty. Check status and correct.<br>– Room temperature is too high. If ambient temperature cannot be reduced in a reasonable time, turn OFF product to prevent shortened life.<br>– Temperature sensor is faulty. Compare temperature reported by peer BCC. |
| | Major Warning | none | Non-critical Temp is >36º C | |
| **Temperature sensors are on the BCC and are independent of power supplies. Investigate temperature warnings immediately, before power supplies sense critical temperature and turn off.** | | | | |
| Temperature is under limit | Critical | none | Critical Temp is <9.5º C | – Room temperature is too low.<br>– Temperature sensor is faulty. Compare temperature reported by peer BCC. |
| | Major Warning | none | Non-critical Temp is <15.5º C | |
| Voltage is over limit | Critical | none | Critical | Power supply is faulty. Check status and correct. |
| | Major Warning | none | Non-critical | |
| Voltage is under limit | Critical | none | Critical | Either power supply is faulty. Check status and correct. |
| | Major Warning | none | Non-critical | |

**Table 52** Disk Enclosure Troubleshooting Table (cont'd)

| Problem Description | HW Event Category | LED State | Status | Probable Cause/Solution |
|---|---|---|---|---|
| Peer BCC status, temperature and voltage are Not Available | Major Warning | none | Both BCCs: Non-critical | Firmware on BCC A and BCC B are different versions. |
| | | none | Not Available | Internal bus is faulty. Contact HP technical support to replace midplane. |

# 7 REMOVAL AND REPLACEMENT



Overview . . . . . 384

Disk Enclosure Modules . . . . . 386

Controller Enclosure Modules . . . . . 396

# Overview

This chapter describes removal and replacement procedures for the disk array hot-swappable modules that are customer replaceable. Hot-swappable modules can be replaced without impacting host interaction with the disk array.

Procedures for replacing the following modules are included in this chapter:

- Disk Enclosure
  - Disk Modules
  - Power Supply Modules
  - Fan Modules
- Disk Array Controller Enclosure
  - Battery Backup Module
  - Controller Fan Module
  - Power Supply Modules
  - Power Supply Fan Module

All modules other than those listed above must be replaced by a trained service representative.

**Note** **What is a “hot-swappable” component?**
The term hot-swappable identifies a component that can be removed and replaced without disrupting host access to the disk array. Host I/Os continue without interruption when replacing a hot-swappable component. Disk array controller modules, disk modules, power supply modules, and fan modules are all hot swappable components.

**Note**

**Is the HP SureStore E Disk Array FC60 customer repairable?**
Although the modular design of the Disk Array FC60 makes it easy to replace failed components, it is recommended that repair of the product be done by a trained service representative. This includes troubleshooting, and removal and replacement of hot-swappable components. Customer repair should be performed only in situations where a representative is not available, or system operation may be impacted if repair is not performed immediately.

The HP SureStore E Disk Array FC60 is fully covered by a warranty from Hewlett-Packard. Additional support services may have also been purchased for the disk array. During the warranty period, or if the product is covered by a service contract, it is recommended that you contact your service representative for all service and support issues.

# Disk Enclosure Modules

This section describes the procedures for replacing the hot swappable modules in the disk enclosure.

## Disk Module or Filler Module

☑ Hot Swappable Component!

This procedure describes how to add or replace disk modules and disk slot filler modules. When adding or replacing disk filler modules use the same procedure, ignoring any steps or information that applies to disk modules only.

***CAUTION*** *- To prevent damage from static electricity, follow standard ESD procedures and avoid touching exposed circuitry.*



*- Do not remove a disk module or filler module until you have a replacement. An empty slot will disrupt the cooling airflow within the enclosure and may cause overheating.*

*- To avoid removing the wrong disk, it is recommended that you identify the disk by flashing its amber Fault LED. Identifying a disk can be done using the management software tools. Removing the wrong disk can cause data to become unavailable if the disk is part of the same LUN as the failed disk.*

*- Never move disks from one array to another or within an array with power on. Moving disks with power on may result in data loss. Always power down the disk array before moving disks. Moving disk should not be confused with replacing a failed disk, which can be done with power on.*

**Note** When a disk module is replaced, the new disk inherits the group properties of the original disk. For example, if you replace a disk that was part of LUN 1, the replacement will also become part of LUN 1. If the disk is a replacement for a global hot spare or an unassigned disk, the replacement will become a global hot spare or an unassigned disk.

A special feature called drive lockout prevents unsupported disk drives from being used in the disk array. If an unsupported disk drive is installed in the disk array, the drive will be failed. Locked out drives will be indicated with a special drive status (0x1B).

## Removing a Disk Module or Filler Module

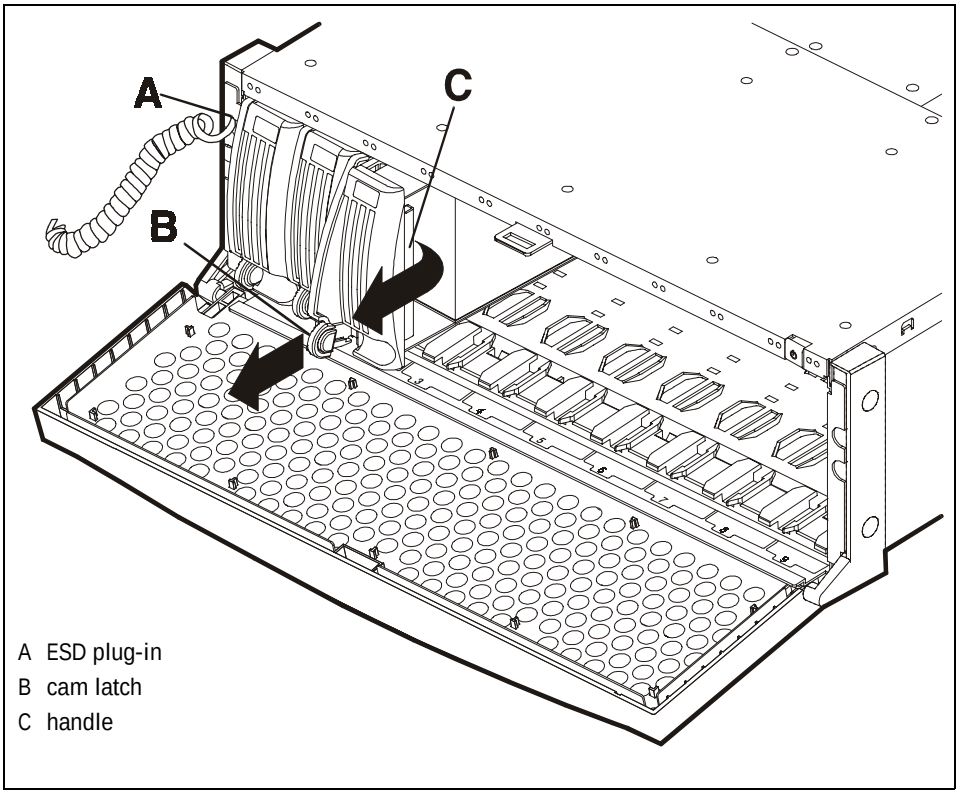
1. Unlock and open the disk enclosure door.
2. If removing a disk module, insert the plug end of your ESD wrist strap in the disk enclosure ESD socket (A in Figure 94).
3. Release the module by squeezing the latch tab (B in Figure 94) and sliding the module part way out of the enclosure.

***CAUTION*** *While the disk is spinning, it is vulnerable to damage. Wait for the disk to stop spinning (about 15 seconds) before fully removing the module from the enclosure.*

**WARNING** **To avoid electric shock hazard, do not touch the backplane or adjacent drive electronics when removing and inserting disks.**

4. Slide the module out of the slot until you can grasp the handle (C). Support the disk module with your other hand around the enclosed side.

**Note** When removing adjacent disk modules, removing them from right to left improves access to successive modules.




A ESD plug-in
B cam latch
C handle

**Figure 94** Disk Module Removal

## Installing a Disk Module or Filler Module

***CAUTION*** *Touching the disk circuit board can cause high energy discharge and damage the disk.*



*Disks modules are fragile and should be handled carefully.*



**Note** If the disk module you are installing has been removed from another Disk Array FC60, you should ensure that the module has a status of Unassigned. This is done by unbinding the LUN the disk module was a part of in the original disk array. See "Moving Disks from One Disk Array to Another" on page 255.

1. Remove the replacement disk from its ESD bag, being careful to grasp the disk by its handle (A in Figure 96).
2. Pull the cam latch (B) away from the disk module.
3. Mark the slot number that the disk is being installed into on the Slot Location Map on the top of the disk module. See Figure 95. This facilitates reinstalling the disk module in the correct slot.

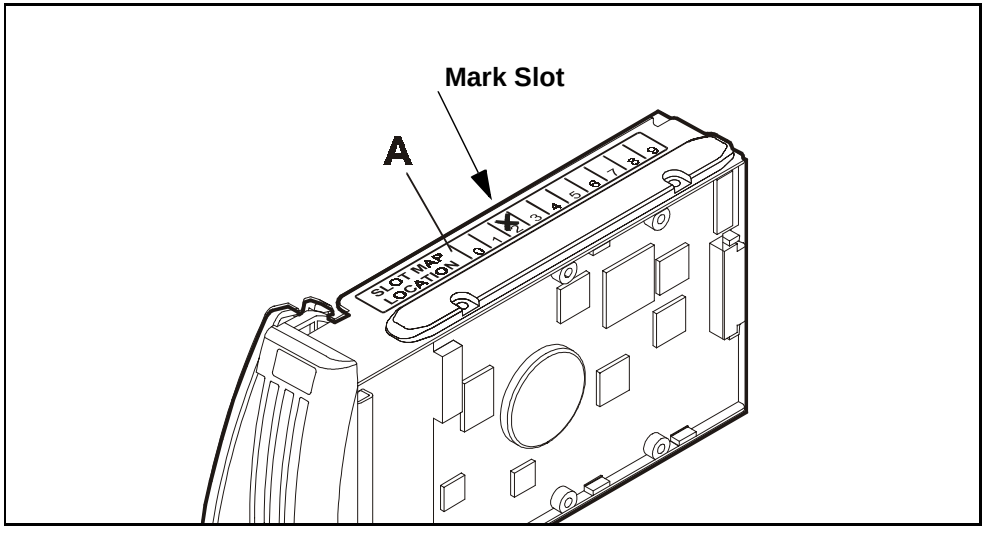



**Figure 95** Disk Slot Identification Label

**Note** Use both hands to hold the disk — one on the handle and the other on the carrier frame.

4. Slide the disk module into the slot with capacity label up (C in Figure 96).

5. Close the cam latch to seat the module firmly into the backplane. An audible click indicates the latch is closed properly.
6. Check the LEDs (D in Figure 96) above the disk module for the following behavior:
   - Both LEDs should turn on briefly.
   - The amber Fault LED should turn off.
   - The green disk Activity LED should blink for a few seconds and then go out. If the host begins to access the disk, the Activity LED will flash.

   If this behavior does not occur, refer to chapter 6, Troubleshooting, for information on identifying and solving the problem.
7. Unplug the ESD strap, and close and lock the disk enclosure door.
8. Disable WCE on the new disk module using the following command:

   **amutil -w <ArrayID>**

   See "Disabling Disk Module Write Cache Enable (WCE)" on page 302 for more information.

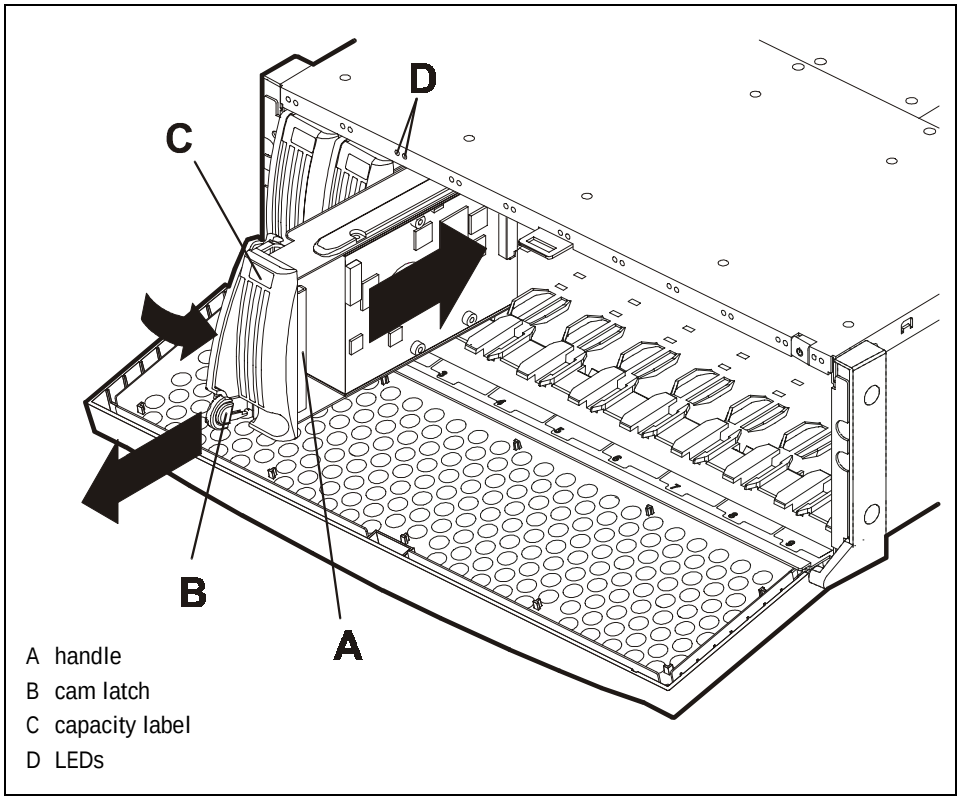



**Figure 96** Disk Module Replacement

## Disk Enclosure Fan Module

☑ Hot Swappable Component!

A failed fan module should be replaced as soon as possible. There are two fan modules in the enclosure. If a fan fails, the remaining fan module will maintain proper cooling. However, if the remaining fan module fails before the defective fan is replaced, the disk enclosure must be shut down to prevent heat damage.

***CAUTION*** *Do not remove a disk fan module until you a replacement. An empty slot will disrupt the cooling airflow within the enclosure and may cause overheating.*

### Removing the Fan Module

1. Loosen the two module locking screws (A in Figure 97).
2. Slide the fan module out of the chassis by the pull tab (B).

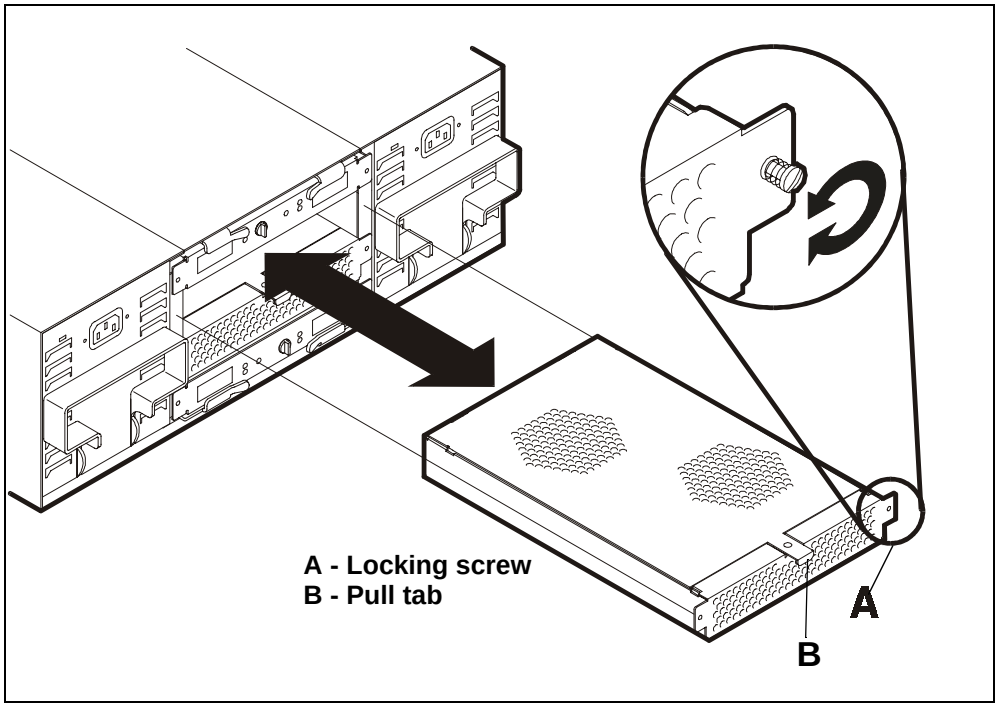



**Figure 97** Disk Enclosure Fan Module Removal and Replacement

## Installing the Fan Module

1. Slide the replacement fan module into the empty slot (C in Figure 97).
2. Tighten the locking screws (A).
3. Check the fan module LED for the following behavior:
   – The Fan Fault LED should flash amber, and then turn green after a few seconds. If the LED does not turn green, refer to "Troubleshooting" on page 359.

## Disk Enclosure Power Supply Module

☑ Hot Swappable Component!

A failed power supply module should be replaced as soon as possible. When one power supply fails, the remaining power supply will maintain the proper operating voltage for the disk enclosure. However, if the remaining power supply fails before the first power supply is replaced, all power will be lost to the disk enclosure.

***CAUTION*** *Do not remove a power supply module until you have a replacement. An empty slot will disrupt the cooling airflow within the enclosure and may cause overheating.*

### Removing the Power Supply Module

1. Disconnect the power cord to the power supply.
2. Loosen the screw (B in Figure 98) from the power supply handle (A).
3. Pull the handle down to disengage the power supply from the backplane.
4. Slide the power supply out of the chassis. Support the power supply with your free hand as it clears the chassis.

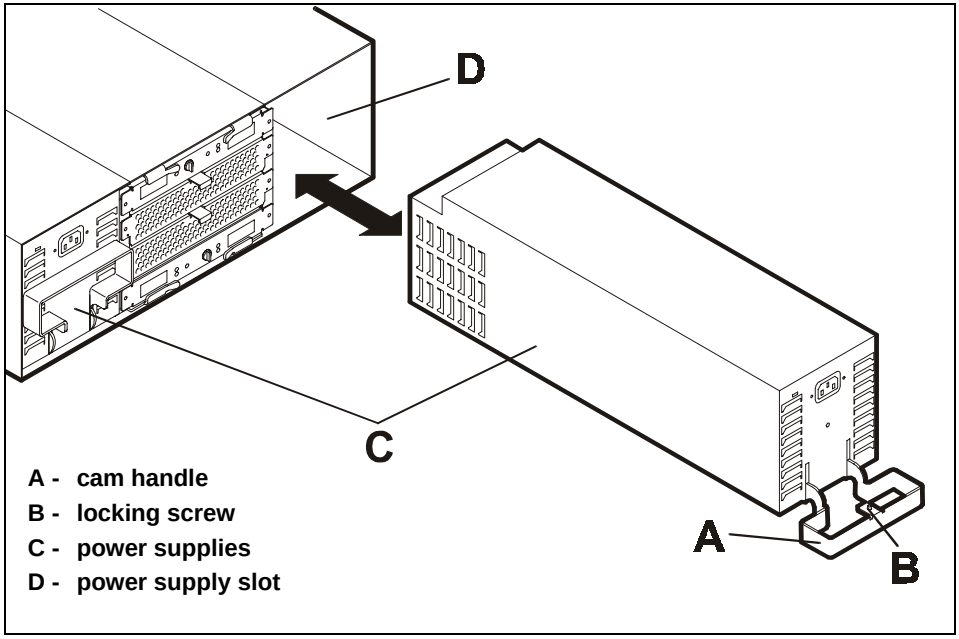



**Figure 98** Disk Enclosure Power Supply Module Removal and Replacement

## Installing the Power Supply Module

1. With the handle down, slide the replacement power supply into the empty slot (D in Figure 98). The supply begins to engage the backplane with 3/8 inch (8 mm) still exposed.
2. Swing the handle upward to seat the power supply into the backplane. The power supply should be flush with the chassis.
3. Tighten the screw (B) in the power supply handle (A ).
4. Plug the power cord into the power supply and electrical source.
5. Check the Power Supply LED. It should then turn green. If the LED remains off or is amber, see chapter 6, Troubleshooting for information on identifying and solving the problem.

# Controller Enclosure Modules

This section provides removal and replacement procedures for the controller enclosure modules, plus the controller enclosure front cover. Most controller modules are hot swappable, however certain restrictions need to be observed for some modules, as identified in these descriptions.

The controller modules, the controller fan module, and the BBU are accessed from the front of the controller enclosure. Access to these modules requires that the front cover be removed. The power supply modules and the power supply fan module are accessed from the rear of the controller enclosure.

***CAUTION*** *Do not remove a module from the controller enclosure until you are ready to install the replacement. An empty slot will disrupt the cooling airflow within the enclosure and may cause overheating.*

## Front Cover Removal/Replacement

☑ Hot Swappable Component!

To gain access to the front of the controller module, the controller fan module, or the battery backup unit (BBU), the front cover must be removed.

### Removing the Front Cover

1. Pull the bottom of the cover out about one inch to release the pins. See Figure 99.
2. Slide the cover down one inch and pull it away from the controller enclosure.

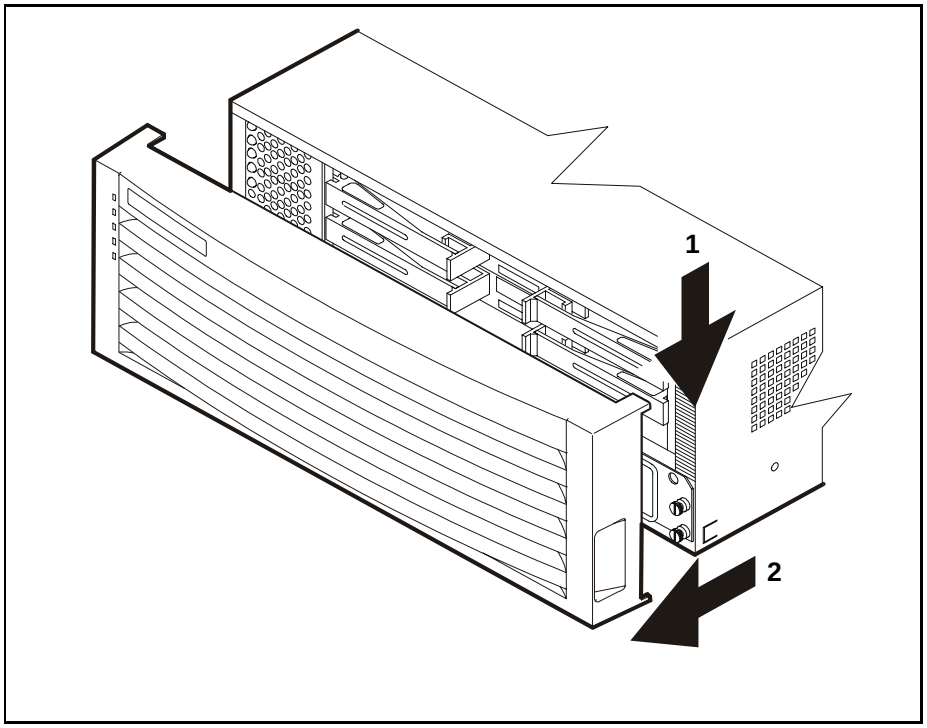



**Figure 99** Controller Enclosure Front Cover Removal and Replacement

### Installing the Front Cover

1. Slide the top edge of the cover up under the lip of the chassis.
2. Push the cover up as far as it will go, then push the bottom in until the pins snap into the mounting holes.

## Controller Fan Module

☑ Hot Swappable Component!

***CAUTION*** *Do not operate the controller enclosure without adequate ventilation and cooling to the controller modules. Operating without proper cooling to the controller modules may damage them. You must shut down the controller enclosure to avoid overheating if you anticipate needing more than 15 minutes to replace the fan module.*

### Removing the Controller Fan Module:

1. Remove the controller enclosure front cover. See "Front Cover Removal/Replacement" on page 397.
2. Loosen the screw on the fan module. See Figure 100.
3. Slide the controller fan module out of the enclosure.

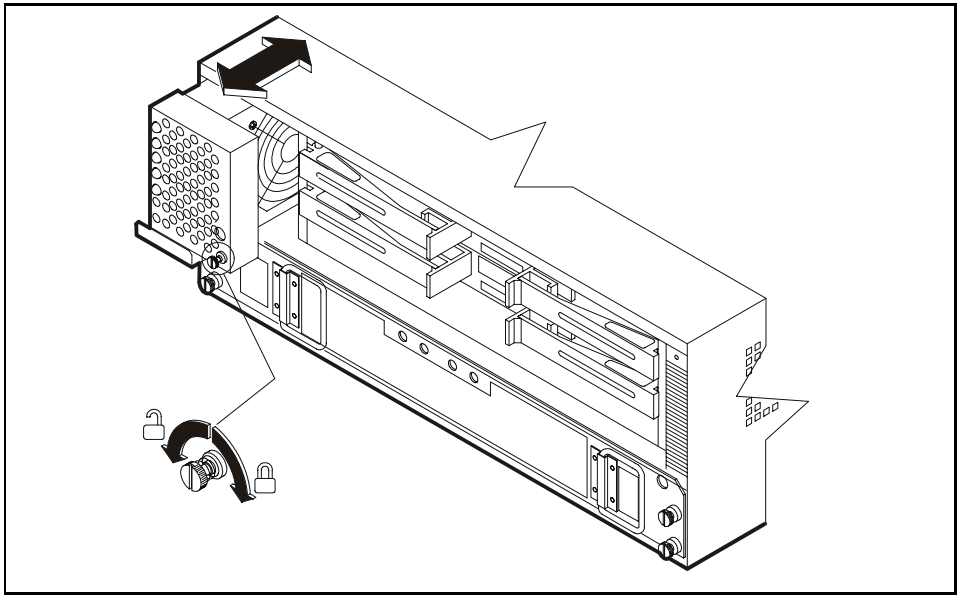


**Figure 100** Controller Fan Module Removal and Replacement

## Installing the Controller Fan Module.

1. Slide the new module into the slot and tighten the screw. The captive screw is spring-loaded and will not tighten unless it is inserted all the way into the chassis. If the screw keeps spinning without tightening, push on the front of the module until it snaps into place, and then tighten the screw.
2. Check the system LEDs on the controller fan module for the following behavior:
   - The green Power LED should be on and the Fan Fault LED should be off.
     If the Fan Fault LED turns on or the Power LED remains off, make sure that the controller enclosure fan module is inserted all the way into the chassis, and the thumbscrew is tight. If the Fan Fault LED remains on, it may indicate a problem with the new module.
3. Install the front cover. See "Front Cover Removal/Replacement" on page 397.

## Battery Backup Unit (BBU) Removal/Replacement

☑ Hot Swappable Component!

**Note** If the Fast Write Cache LED is on when the BBU is removed from the enclosure (or if the BBU fails), write caching will be disabled and the write cache data will be written to disk. However, if a power outage occurs prior to completing the cache write to disk, data may be lost. Therefore, make sure the Fast Write Cache LED is off before replacing the BBU.

### Removing the BBU

1. Remove the controller enclosure front cover. See "Front Cover Removal/Replacement" on page 397.
2. Loosen the four screws (two on each side) on the BBU. See Figure 101.

***CAUTION*** *The BBU weighs approximately 24 pounds so be prepared to support its weight when you pull it out of the chassis. The BBU is not as deep as other modules and disengages the enclosure sooner than you may expect.*

3. Grasp the two handles and slide the module out about two inches.
4. Grasp both sides of the module with both hands and slide it out of the enclosure.

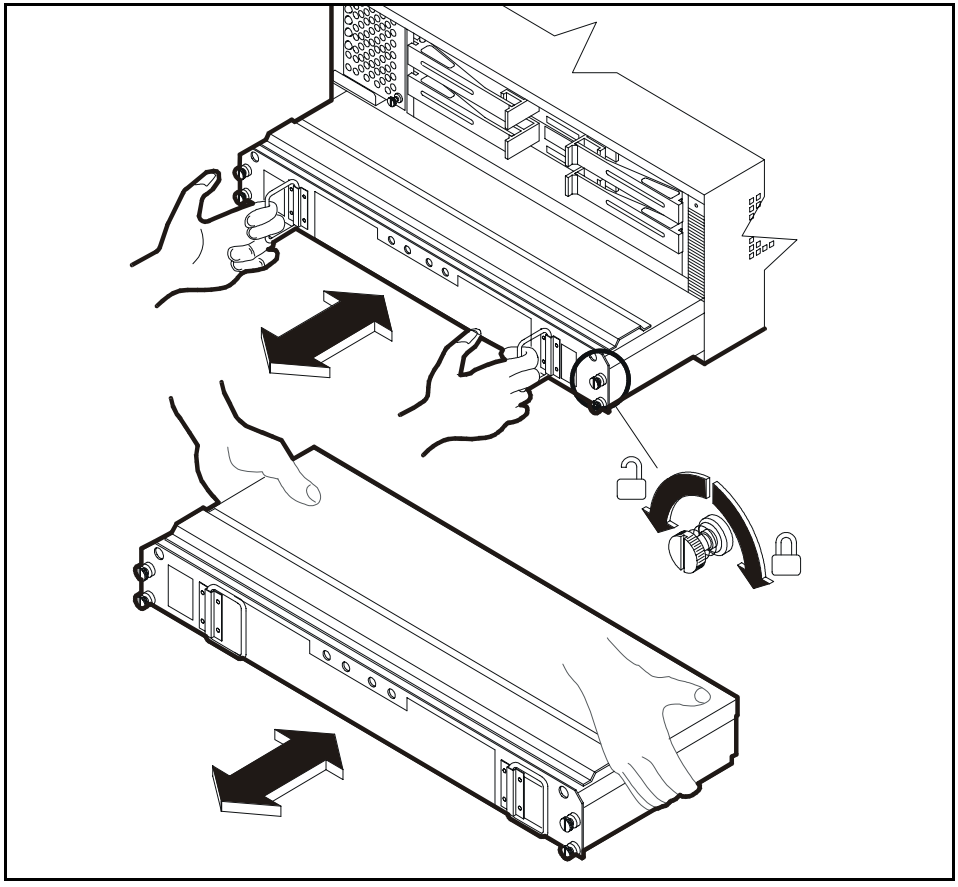


**Figure 101** BBU Removal and Replacement

## Installing the BBU

1. Unpack the new BBU. Save the shipping material for transporting the used BBU to the disposal facility.
2. Fill in the following information on the “Battery Support Information” label on the front of the battery. See Figure 102.
   a. Record the current date on the blank line next to “Date of Installation.”
   b. Record the expiration date (two years from the current date) on the line next to “Replacement Date.”
3. Slide the new BBU into the slot and tighten the screws. See Figure 101.
4. Install the controller enclosure front cover. See "Front Cover Removal/Replacement" on page 397.
5. Allow the system to run up to 7 hours to properly charge the BBU. The Full Charge LEDs will flash while the BBU is charging. When properly charged, both Full Charge LEDs on the front of the BBU will remain on.

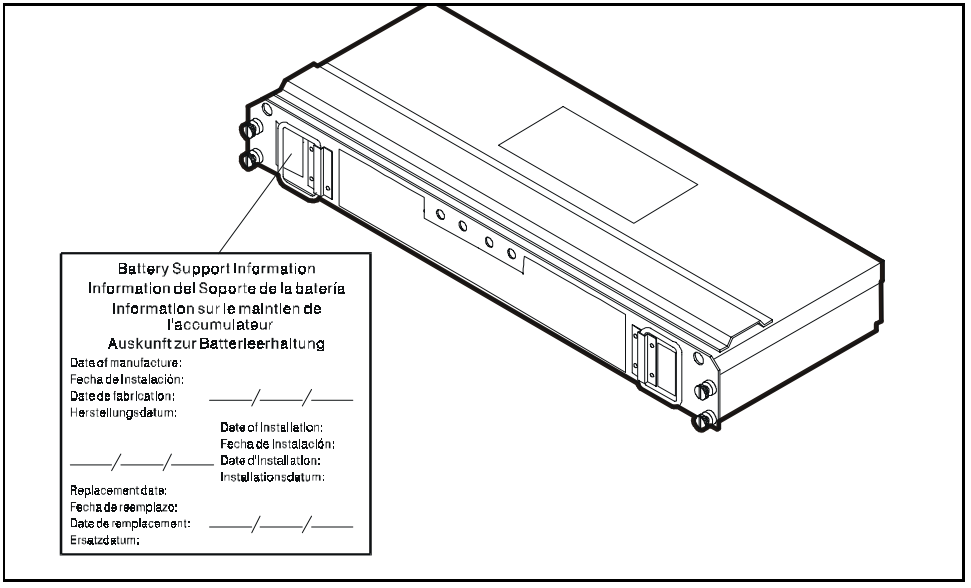



**Figure 102** Battery Information Label

6. Dispose of the old BBU.

**Note** Dispose of the used BBU according to local and federal regulations, which may include hazardous material handling procedures.



## Power Supply Fan Module Removal/Replacement

☑ Hot Swappable Component!

***CAUTION*** *Do not operate the enclosure without adequate ventilation and cooling to the power supplies. Operating the power supplies without proper cooling may damage their circuitry. You must shut down the controller enclosure to avoid overheating if you take more than 15 minutes to remove the Controller Fan Module and insert the new one.*

### Removing the Power Supply Fan Module

1. Gain access to the rear of the disk array controller enclosure.
2. Lift up on the ring/latch on the module to release the catch. See Figure 103.
3. Slide the power supply fan module from the enclosure.

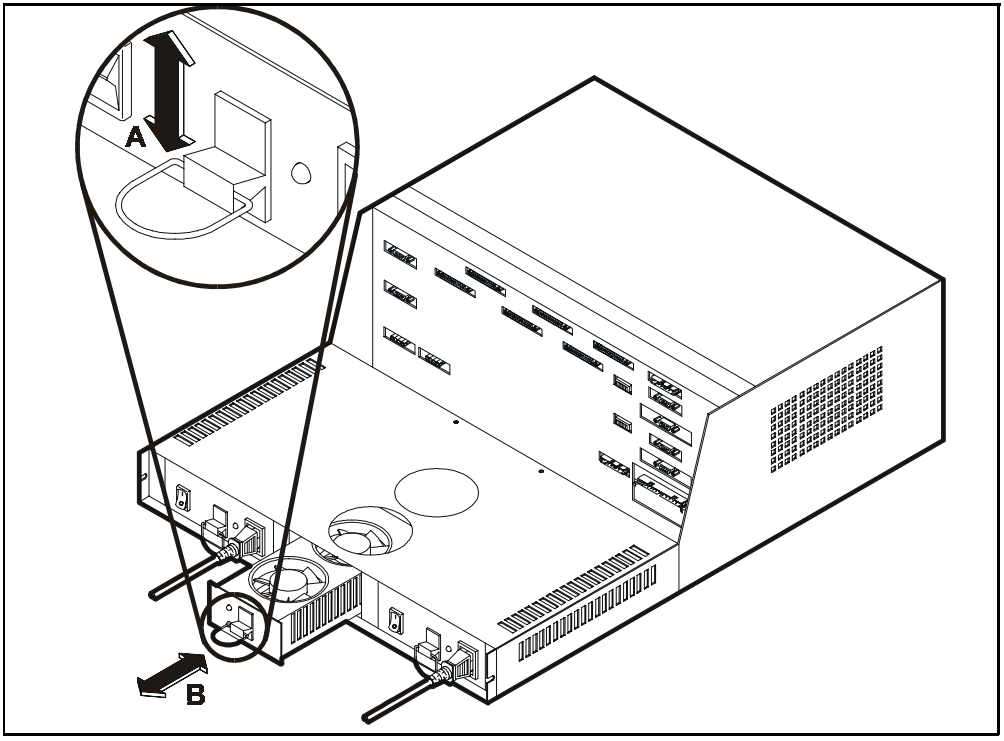



**Figure 103** Power Supply Fan Module Removal and Replacement

### Installing the Power Supply Fan Module

1. Slide the power supply fan module into the enclosure. The latch will snap down when the module is seated properly. If the latch remains up, lift up on the ring/latch and push in on the module until it snaps into place.
2. Check the module LEDs for the following behavior:
   - The green Fan Power LED should be on and the amber Fan Fault LEDs should be off. If the Power LED remains off or the Fault LED turns on, make sure the power supply fan module is inserted all the way into the chassis and locked in place. If the Fault LED remains on, the new power supply fan module may be malfunctioning, refer to Chapter 6, Troubleshooting.

## Power Supply Module Removal/Replacement

☑ Hot Swappable Component!

A power supply should be replaced as quickly as possible to avoid the possibility of the remaining supply failing and shutting down the disk array.

### Removing the Power Supply Module

1. Turn off the power switch and unplug the power cord from the failed power supply module. See Figure 104.
2. Lift up on the pull ring to release the latch. See Figure 105.
3. Slide the supply out of the enclosure.

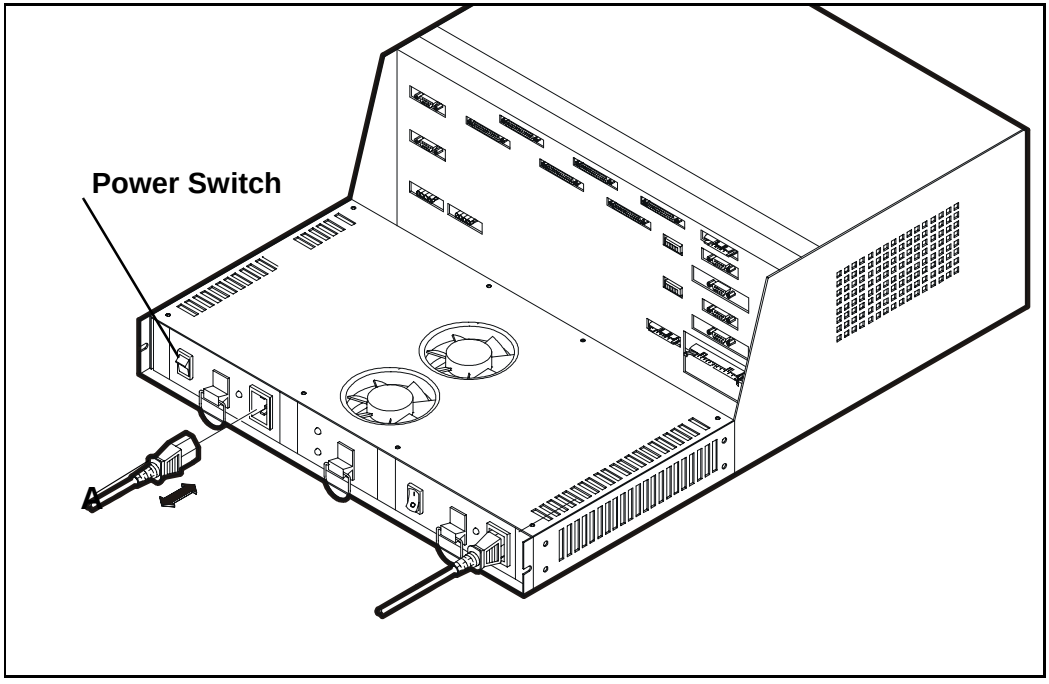



**Figure 104** Removing Power from the Controller Enclosure

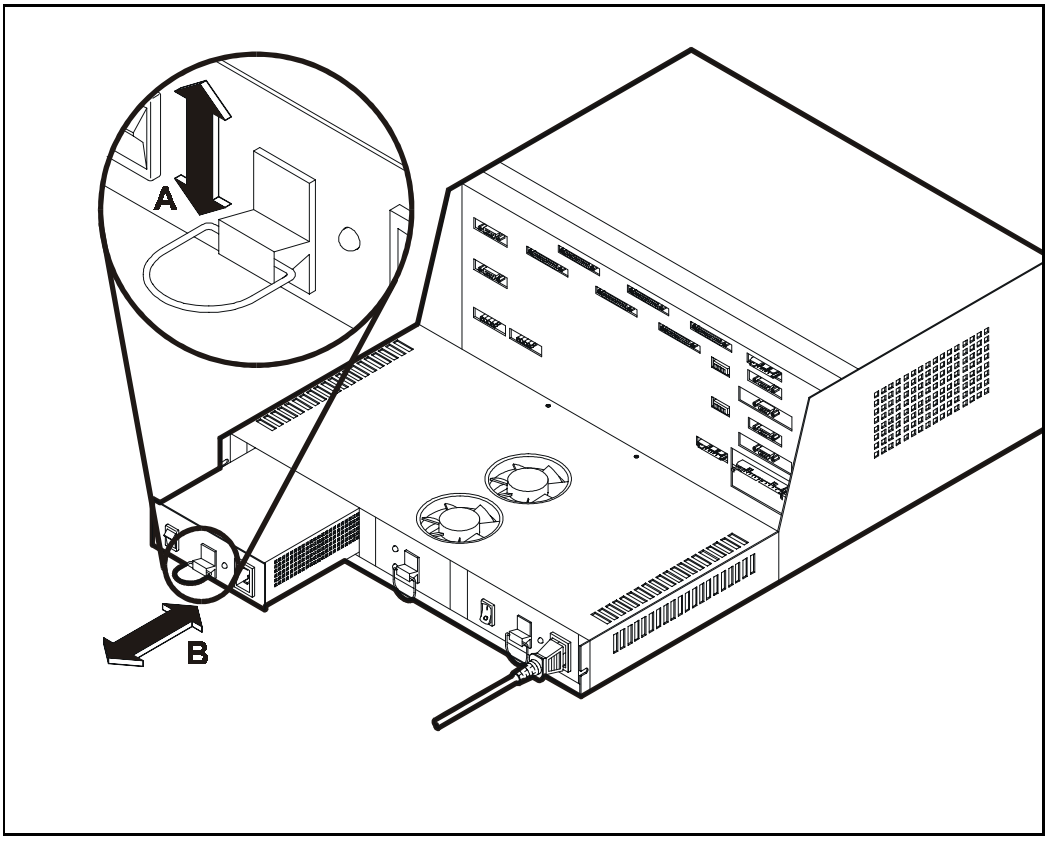



**Figure 105** Power Supply Module Removal and Replacement

### Installing the Power Supply Module

1. Slide the supply into the slot until it is fully seated and the latch snaps into place.
2. Plug in the power cord and turn on the power. See Figure 104.
3. Check the power supply module LED for the following behavior:
   – The Power LED should go on. Once the power supply in installed and operating, there may be a delay of up to several minutes before the Power Fault LED goes off.

## SCSI Cables

Replacing SCSI cables requires that the disk enclosure be shut down. Shutting down the enclosure will degrade the performance of the array during the replacement. When the replacement is completed and the disk enclosure is powered up, the array will perform a rebuild (since I/O has occurred to the array while the disk enclosure was powered off). Array performance will be reduced until the rebuild is complete

To replace a SCSI cable, complete the following steps:

1. Open the disk enclosure front door.
2. Power the disk enclosure down by setting the power switch to off.

***CAUTION*** *When replacing a SCSI cable, it must be disconnected from both the controller enclosure and the BCC module. Disconnect the cable from the controller enclosure first, then from the failed BCC. DO NOT disk connect the SCSI cable from the BCC first or errors may occur when the enclosure is powered on.*

3. Remove the SCSI cable from the controller enclosure connector.
4. Remove the SCSI cable from the BCC
5. Connect the SCSI cable to the BCC connector.
6. Connect the other end of the SCSI cable to the controller enclosure SCSI connector.
7. Disconnect all disk modules from the backplane connectors by releasing the locking levers and pulling module out, about one inch. This step is required to ensure the disk modules return to full operation when the enclosure is powered back up.
8. Power up the disk enclosure.
9. Reseat the disk modules one at a time:
    a. Push the disk module in slot 0 into slot until it snaps into place.
    b. Observe the Disk Fault LED, it should come on. When the Disk Fault LED goes out, insert the next disk module.
    c. Repeat the above steps until all disk modules are re-seated.

Once the disk enclosure is powered up, check the status of the disk modules using one of the software management tools. Initially the disk modules status will be either “write failed” or “no_response.” Eventually, all the disk modules should return to “replaced” status. Once this occurs, the disk array will perform a rebuild (a result of I/O occurring during the period the enclosure was powered off). The disk array will operate at reduced performance until the rebuild is completed.

# 8 REFERENCE / LEGAL / REGULATORY


Models and Options . . . . . 412

PDU/PDRU Products . . . . . 417

Replaceable Parts . . . . . 418

A5277A/AZ Controller Enclosure Specifications . . . . . 420

A5294A/AZ Disk Enclosure Specifications . . . . . 424

Warranty and License Information . . . . . 428

Regulatory Compliance . . . . . 434

FCC Statements (USA Only) . . . . . 435

# System Requirements

## Host Systems

### HP-UX

**Table 53 Supported HP-UX Host Platform Information**

| Supported Host | Boot Support on HP-UX 11.x? | Fibre Channel I/O Adapter |
|---|---|---|
| K-class | Yes | A3404A |
| V-class | Yes | A5158A, A3740A |
| L-class | Yes | A5158A, A3740A |
| D-class | Yes | A3591B |
| N-class | Yes | A5158A, A3740A |
| R-class | Yes | A3591B |
| T-class | No | A3636A |
| C-class | No | A5158A on HP-UX 11.x<br>A3740A on HP-UX 10.20 |
| A4xx-A5xx class | Yes | A5158A |

### Windows NT 4.0 and Windows 2000

Any host running Windows NT 4.0 or Windows 2000.

## Supported Operating Systems

- HP-UX 11.0
- HP-UX 10.20
- Windows NT 4.0 (with Service Pack 4, 5, or 6)

- Windows 2000

## Fibre Channel Host Adapters

### HP-UX

- K-class: A3404A (assy number J2389-60001), 1063 Mbps, short-wave, non-OFC
- D- & R--class: A3591A (number A3395-60001), 1063 Mbps, short-wave, non-OFC
- T-600 class: A3636A (number A3329-60107), 1063 Mbps, short-wave, non-OFC
- V-class: A3740A (number A3740-60001), 1063 Mbps, short-wave, non-OFC

### Windows NT 4.0 and /2000

See the *HP Storage Manager 60 User's Guide* for a list of supported host adapters.

# Models and Options

The HP SureStore E Disk Array FC60 consists of two products: the A5277A/AZ controller enclosure and the A5294A/AZ SureStore E Disk System SC10, or disk enclosure. Each of these products have their own options as indicated below.

## A5277A/AZ Controller Enclosure Models and Options

- **A5277A** is a field racked controller enclosure integrated by qualified service -trained personnel. This model can be ordered with up to six A5294A. disk enclosures.
- **A5277AZ** is a factory-racked controller enclosure integrated into the specified rack. This model must be ordered with from one to six A5294AZ disk enclosures.

Both models include the following components (except as indicated):

- Rackmount enclosure
- Two controller modules with 256 Mbytes of cache (single controller available)
- Two power supplies
- Two fan modules
- One battery backup unit (BBU)
- Two power cords (w/ ferrite bead)
- Mounting rail kits for HP cabinets
- Terminators for unused SCSI ports
- Two 1/2 EIA Filler Panels (A5277A only)
- User documentation

**Table 54** A5277A/AZ Product Options

| Option | Description |
|---|---|
| | **Controller Options (must order one option)** |
| 203 | Single controller module with 256 Mbyte cache, one Media Interface Adaptor, and one controller slot filler module. Configured with HP-UX firmware. |
| 204 | Two controller modules with 256 Mbyte cache and two Media Interface Adaptors. Configured with HP-UX firmware. |
| 205 | Two controller modules with 256 Mbyte cache and two Media Interface Adaptors. Configured with Windows NT/2000 firmware and Windows NT NVSRAM. Includes Storahe |
| 304 | Two controller modules with 512 Mbyte cache and two Media Interface Adaptors. Configured with HP-UX firmware. |
| 305 | Two controller modules with 512 Mbyte cache and two Media Interface Adaptors. Configured with Windows NT firmware. |
| | **Host Connect Cable Options** |
| 0Z4 | 2-meter Fibre Channel cable |
| AFY | 16-meter Fibre Channel cable |
| 0Z5 | 50-meter Fibre Channel cable |
| 0Z6 | 100-meter Fibre Channel cable |
| 701 | Replace 1.5M VHDCI SCSI cable with 5M VHDCI SCSI cable |
| | **Operating System Support Option (No longer required)** |
| UM4 | Support for HP-UX 11.0 |

## A5294A/AZ Disk Enclosure SC10 Models and Options

Order the following product and options as required. Enter the following product and options as sub-items to the A5277A and A5277AZ products above.

- A**5294A disk enclosure** is a field racked Sure Store E Disk System SC10 integrated by a service-trained engineer. This product may be ordered in conjunction with A5277A controller enclosure. To order a disk system SC10 without integration into an array, order A5272A.
- **A5294AZ** Factory-racked Sure Store E Disk System SC10 (disk enclosure) integrated into the specified rack with the A5227AZ controller. This product may only be ordered in conjunction with A5277AZ controller enclosure. To order a disk enclosure SC10 without integration into an array, order A5272AZ

Both models include the following components (except as indicated):

- SC10 Rack-mount enclosure (accommodates 10, 1.6" or 1," disk modules)
- Two power supplies
- Two fan modules
- Two power cords
- Two bus controller cards (BCC; with enclosure monitoring)
- Mounting rails for HP racks
- Two SCSI bus terminators
- User documentation
- 1/2 EIA filler panel (A5294AZ as required for proper rack appearance)
- One or two 2.0 meter SCSI cables - The number of cables included depends on the number of A5294A/AZ disk enclosures ordered:

| Number of A5294A/AZ per A5277A/AZ | Number of SCSI cables per A5294A/AZ |
|---|---|
| 1,2, or 3 | 2 |
| 4, 5, or 6 | 1 |

**Table 55** A5294A Custom Cabling Option

| Option | Description |
|---|---|
| 701 | Delete one 2m cable included in A5294A product and add one 5m VHDCI SCSI cable for connection of A5277A to A5294A in a different rack |

**Table 56** A5294A/AZ Storage Capacity Options

| Option | Description |
|---|---|
| Note: All disk enclosures ordered with a single A5277A/AZ must have identical Storage Capacity Options. | |
| 104 | 4, 9-Gbyte 10K rpm disk drives |
| 108 | 8, 9-Gbyte 10K rpm disk drives |
| 110 | 10, 9-Gbyte 10K rpm disk drives |
| 204 | 4, 18-Gbyte 10K rpm disk drives |
| 208 | 8, 18-Gbyte 10K rpm disk drives |
| 210 | 10, 18-Gbyte 10K rpm disk drives |
| 304 | 4, 36-Gbyte 10K rpm disk drives |
| 308 | 8, 36-Gbyte 10K rpm disk drives |
| 310 | 10, 36-Gbyte 10K rpm disk drives |
| 504 | 4, 18-Gbyte 15K rpm disk drives |
| 508 | 8, 18-Gbyte 15K rpm disk drives |
| 510 | 10, 18-Gbyte 15K rpm disk drives |

**Table 57** Supporting Software (CD ROMS)

| Product | Description |
|---|---|
| B6191A | IPR Distribution CD-ROM |
| A5628A | *HP Storage Manager 60* |

## Disk Array FC60 Upgrade and Add-On Products

Order the following parts to expand or reconfigure your original purchase:

**Table 58** Upgrade Products

| Order No. | Description |
|---|---|
| A5276A | 9.1-Gbyte disk drive module 10K rpm Ultra 2 LVD |
| A5282A | 18.2-Gbyte disk drive module 10K rpm Ultra 2 LVD |
| A5633A | 18.2-Gbyte disk drive module 15K rpm Ultra 2 LVD |
| A5595A | 36.4-Gbyte disk drive module 10K rpm Ultra 2 LVD |
| A5622A | 73.4-Gbyte disk drive module 10K rpm Ultra 2 LVD |
| A5278A | Add-on controller, includes Media Interface Adaptor (MIA)<br>Must also order Option 002 for 256 Mbytes of cache |
| 002 | 256 Mbyte cache option for A5278A |
| A5279A | Controller cache 512-Mbyte upgrade kit. Includes two 256-Mbyte DIMMs |
| A5306A | 2 m, VHDCI to VHDCI SCSI Cable, M/M |
| A5307A | 5 m, VHDCI to VHDCI SCSI Cable, M/M |
| A5308A | 10 m, VHDCI to VHDCI SCSI Cable, M/M |
| A5250A | Rail kit for HP legacy cabinets: C2785A, C2786A, C2787A, A1896A, and A1897A |
| A5251A | Rail kit for HP Rack System/E: HP A4900A, A4901A, A4902A, J1500A, J1501A, J1502A |
| A5672A | Rail kit for Rittal 9000 racks |
| A3583A | 2-meter Fibre Channel cable |
| A3531A | 16-meter Fibre Channel cable |
| A3735A | 50-meter Fibre Channel cable |
| A3736A | 100-meter Fibre Channel cable |
| A5296A | VHDCI SCSI terminator (LVD/SE) |
| A5649A | Storage Partition Support (up to 8 partitions) |

# PDU/PDRU Products

Hewlett-Packard offers the following PDUs and PDRUs, with US and international power options, for meeting electrical requirements:

**Table 59** PDU/PDRU Products

| Order No. | Description |
|---|---|
| | Supported on Original Racks |
| E7676A | 19 inch, 100-240 V, 16 Amp, 1 C20 inlet, 10 C20 outlets |
| E7671A | 19 inch, 100-240 V, 16 Amp, 1 C20 inlet, 2 C19 & 6 C13 outlets |
| E7674A | 19 inch, 100-240 V, 16 Amp, 1 C20 inlet, 1 C19 & 7 C13 outlets |
| E7679A | 19 inch, 100-127 V, 16 Amp, 2 C20 inlets, 2 C19 outlets, switch accessory |
| E7680A | 19 inch, 200-240 V, 16 Amp, 2 C20 inlets, 2 C19 outlets, switch accessory |
| E7681A | 19 inch, 200-240 V, 30 Amp, L6-30P, 2 C19 & 8 C13 outlets, switch accessory |
| E7682A | 19 inch, 200-240 V, 30 Amp, IEC-309, 2 C19 & 2 C13 outlets, switch accessory |
| E4452A | 36 inch, 220 V, 16 Amp, L6-20P, 6 C-13 outlets |
| E4453A | 36 inch, 220 V, 16 Amp, no plug, 6 C-13 outlets |
| E5933A | 36-inch, 110-220 V, 16 Amp, UPS, IEC-320, 10 C-13 outlets |
| E4456A/B | 60 inch, 220 V, 16 Amp, L6-20P, 10 C-13 outlets |
| E4457A/B | 60 inch, 220 V, 16 Amp, no plug, 10 C-13 outlets |
| E5930A | 60 inch, 110-220 V, 16 Amp, UPS, IEC-320, 10 C-13 outlets |
| E5931A | 60 inch, 220 V, 16 Amp, UPS, LP-30P, 10 C-13 outlets |
| E5932A | 60 inch, 220 V, 16 Amp, UPS, no plug, 10 C-13 outlets |
| E7677A | Switch panel accessory for PRU |
| E7678A | Switch control jumper cord for PRU |

# Replaceable Parts

## A5277A/AZ Controller Enclosure Replaceable Parts

Table 60 Controller Enclosure Replaceable Parts

| Part Number | Field Replaceable Units | Exchange Part # |
|---|---|---|
| A5278-60001 | HP-UX[1] Controller Module (5v model[2]) w/32 MB SIMM (no cache DIMMs) This part has been replaced by the A5278-60006. | A5278-69001 |
| A5278-60006 | HP-UX[1] Controller Module (3.5v model[2]) w/32 MB SIMM (no cache DIMMs) | A5278-69006 |
| A5635-60002 | Windows NT[1] Controller Module w/32 MB SIMM (no cache DIMMs) Includes Windows NT NVSRAM settings | A5635-69002 |
| A5278-60002 | 128 MB DIMM | n/a |
| 5065-5201 | 256 MB DiMM | n/a |
| A5278-60004 | 16MB SIMM Module | n/a |
| A5277-60009 | Battery Backup Module | A5277-69009 |
| A5277-60003 | Controller Fan Module | n/a |

[1]Controller modules have operating-specific firmware installed and are not interchangeable. The HP-UX 3.5v controller (A5278-60006) can be upgraded to a Windows NT/2000 controller (A5635-60002) using the HP Storage Manager 60-NT software (A5628A). The HP-UX 5v controller (A5278-60001) cannot be upgraded for use on Windows NT/2000.

[2]The 5v HP-UX controller and the 3.5v HP-UX controller are interchangeable. The 5v model will be available until supplies are exhausted and then replaced by the 3.5v model.

Table 60 Controller Enclosure Replaceable Parts (cont'd)

| Part Number | Field Replaceable Units | Exchange Part # |
|---|---|---|
| A5277-60004 | Power Supply Modules | n/a |
| A5277-60002 | Power Supply Fan Module | n/a |
| A5277-60001 | Front Door Assembly | n/a |
| 5021-1121 | Terminator, SCSI, 68 pin, LVD | n/a |
| 5064-2464 | Media Interface Adapter (MIA) | n/a |

## A5294A/AZ Disk Enclosure Replaceable Parts

**Table 61** Disk Enclosure Replaceable Parts

| Replacement Part Order No. | Part Description | Exchange Part Order No. |
|---|---|---|
| 8120-6514 | Power cord | n/a |
| A5236-60003 | Fan | n/a |
| A5236-60023 | Power supply | A5236-69023 |
| A5272-67004 | Bus Control Card (BCC) | A5272-69004 |
| A5236-60009 | Disk filler assembly | n/a |
| A5276-67001 | 9.1-Gbyte disk drive module 10K rpm Ultra 2 LVD | A5276-69001 |
| A5282-67001 | 18.2-Gbyte disk drive module 10K rpm Ultra 2 LVD | A5282-69001 |
| A5633-67001 | 18.2-Gbyte disk drive module 15K rpm Ultra 2 LVD | A5633-69001 |
| A5595-67001 | 36.4-Gbyte disk drive module 10K rpm Ultra 2 LVD | A5595-69001 |
| A5622-67001 | 73.4-Gbyte disk drive module 10K rpm Ultra 2 LVD | A5622-69001 |

# A5277A/AZ Controller Enclosure Specifications

## Dimensions:

| Height | Width | Depth |
|---|---|---|
| 6.75 inches (17.1 cm) | 17.5 inches (44.5 cm) | 24 inches (61 cm) |

## Weight:

| Component | Weight of Each (lbs) | Quantity | Subtotal (lbs) |
|---|---|---|---|
| Controller modules | 6.6 | 2 | 13.2 |
| Controller Fan | 1.9 | 1 | 1.9 |
| Battery | 21.4 | 1 | 21.4 |
| Power Supply | 3.3 | 2 | 6.6 |
| Power Supply Fan | 1.5 | 1 | 1.5 |
| Front Cover | 2 | 1 | 2 |
| Chassis | 31.6 | 1 | 31.6 |
| Controller Enclosure - Total | | | 78 lbs |

## AC Power:

### AC Voltage and Frequency:

- 120 VAC (100 - 127 VAC), 50 to 60 Hz single phase
- 230 VAC (220 - 240 VAC), 50 to 60 Hz single phase
- Auto-ranging

### Current:

| Voltage | Typical Operating Current | Maximum Operating Current | In-Rush Current |
|---|---|---|---|
| 100 - 127 VAC | 1.5 A | 2.3 A | 21.7 A |
| 220 - 240 VAC | 0.8 A | 1.2 A | 42.9 A |

### Power Consumption:

| Incoming Voltage AC RMS | Typical Power Consumption |
|---|---|
| 100- 127 VAC | 180 watts |
| 200 - 240 VAC | 180 watts |

## Heat Output:

- 615 BTU/hr.

## Environmental Specifications

**Note** The HP SureStore E Disk Array FC60 has been tested for proper operation in supported Hewlett-Packard cabinets. If the disk array is installed in an untested rack configuration, care must be taken to ensure that all necessary environmental requirements are met. This includes power, airflow, temperature, and humidity. Failure to meet the required operating specifications may result in product failure.

### Operating Environment

For continuous, trouble-free operation, the high availability disk array should not be operated at its maximum environmental limits for extended periods of time. The recommended operating range provides a less stressful operating environment. Operating within the recommended operating range ensures maximum reliability for the disk array.

The following specifications were type-tested under controlled conditions. Hewlett-Packard maintains an active program of auditing production products to ensure these specifications remain true when products are again tested under the same conditions. The limits of these specifications do not represent the optimum for long, trouble-free operation and are specifically not recommended for maximum customer satisfaction. The recommended conditions are stated separately where appropriate.

- Operating temperature: 5º C to 40º C (41º F to 104º F)
  **Recommended**: 20º C to 25.5º C (68º F to 78º F)
- Maximum gradient: 20º C per hour (36º F per hour)
- Relative humidity: 10% to 80% RH at 28º C (wet bulb)
  **Recommended**: 30% to 50%
- Altitude: 3048 m (0 - 10,000 ft)

**Note** For continuous, trouble-free operation, the product should NOT be operated at its maximum environmental limits for extended periods of time. Operating within the recommended operating range, a less stressful operating environment, ensures maximum reliability.

### Non-operating Environmental (shipping and storage):

- Temperature: -40º C to 70º C (-40º F to 158º F)
- Maximum gradient: 20º C per hour (68º F per hour)
- Relative humidity: 10% to 90% RH @ 28º C (wet bulb)
- Altitude: 4572 m (0 - 15,000 ft)

## Acoustics

- Meets or exceeds all known international acoustics specifications for computing environments.

### Battery Back-Up Module Disposal

The BBU contains batteries that may be harmful to the environment if not disposed of properly. Please follow the local, state or country regulations regarding the proper disposal of these batteries.

# A5294A/AZ Disk Enclosure Specifications

## Dimensions:

| Height | Width | Depth |
|---|---|---|
| 5.91 in. (15.0 cm) | 18.9 in. (48.0 cm) | 27.2 in. (69.1 cm) |

## Weight:

| Component | Weight of Each (lbs) | Quantity | Subtotal (lbs) |
|---|---|---|---|
| Disk Drive (HH) | 2.8 | 10 | 28 |
| Fan | 3 .3 | 2 | 7 |
| Power Supply | 10.6 | 2 | 22 |
| BCC | 4.5 | 2 | 9 |
| Midplane-Mezzanine | 6 | 1 | 6 |
| Door | 2 | 1 | 2 |
| Chassis | 35 | 1 | 35 |
| Total, Approx. | | | 110 lbs. (50 kg.) |

## AC Power:

### AC Voltage and Frequency:

- 100 - 127 VAC,  50 to 60 Hz single phase
- 220 - 240 VAC,  50 to 60 Hz single phase:

### Current:

| Voltage | Typical Current | Maximum Current |
|---|---|---|
| 100 - 127 VAC | 4.8 a | 6.5 a |
| 220 - 240 VAC | 2.4 a | 3.2 a |
| 50-60 Hz , single phase, power factor corrected | | |

### Power Consumption:

| Incoming Voltage AC RMS | Typical Power Consumption |
|---|---|
| 100- 127 VAC | 347 watts |
| 200 - 240 VAC | 347 watts |

## DC Power Output:

- Disk: +5 V and +12 V from power supply
- BCC: +5 V and +12 V from power supply, 3.3 V self-generated
- Fan:  +12 V from power supply

## Heat Output:

- 2200 BTU/hr.

## Environmental Specifications

**Note** The HP SureStore E Disk Array FC60 has been tested for proper operation in supported Hewlett-Packard cabinets. If the disk array is installed in an untested rack configuration, care must be taken to ensure that all necessary environmental requirements are met. This includes power, airflow, temperature, and humidity. Failure to meet the required operating specifications may result in product failure.

### Operating Environment

For continuous, trouble-free operation, the high availability disk array should not be operated at its maximum environmental limits for extended periods of time. The recommended operating range provides a less stressful operating environment. Operating within the recommended operating range ensures maximum reliability for the disk array.

The following specifications were type-tested under controlled conditions. Hewlett-Packard maintains an active program of auditing production products to ensure these specifications remain true when products are again tested under the same conditions. The limits of these specifications do not represent the optimum for long, trouble-free operation and are specifically not recommended for maximum customer satisfaction. The recommended conditions are stated separately where appropriate.

### Temperature and Altitude:

- Operating temperature: 5º C to 40º C (50º F to 104º F)
  **Recommended**: 20º C to 25.5º C (68º F to 78º F)
- Maximum gradient: 20º C per hour (36º F per hour)
- Relative humidity: 20% to 80% noncondensing, max. wetbulb at 26º C
  **Recommended**: 30% to 50% noncondensing
- Altitude: 3000 m (10,000 ft)

**Note** For continuous, trouble-free operation, the disk enclosure should NOT be operated at its maximum environmental limits for extended periods of time. Operating within the recommended operating range, a less stressful operating environment, ensures maximum reliability.

The environmental limits in a nonoperating state (shipping and storage) are wider:

- Temperature: -40º C to 70º C (-40º F to 158º F)
- Maximum gradient: 24º C per hour (43.2º F per hour)
- Relative humidity: 15% to 90% noncondensing
- Altitude: 4600 m (15,000 ft)

## Acoustics

- Sound power: 6.4 Bels
- Sound pressure at operator's position: 56.2 dB(A)

# Warranty and License Information

## Hewlett-Packard Hardware Limited Warranty

HP warrants to you, the end-user Customer, that HP SureStore E Disk Array FC60 hardware components and supplies will be free from defects in material and workmanship under normal use after the date of purchase for *three years*. If HP or Authorized Reseller receives notice of such defects during the warranty period, HP or Authorized Reseller will, at its option, either repair or replace products that prove to be defective. Replacement parts may be new or equivalent in performance to new.

Should HP or Authorized Reseller be unable to repair or replace the hardware or accessory within a reasonable amount of time, Customer's alternate remedy will be a refund of the purchase price upon return of the HP SureStore E Disk Array FC60.

**Replacement Parts Warranty**

HP replacement parts assume the remaining warranty of the parts they replace. Warranty life of a part is not extended by means of replacement.

**Items Not Covered**

Your HP SureStore E Disk Array FC60 warranty does not cover the following:

- Products purchased from anyone other than HP or an authorized HP reseller
- Non-HP products installed by unauthorized entities
- Consumables, such as batteries
- Customer-installed third-party software
- Routine cleaning, or normal cosmetic and mechanical wear
- Damage caused by misuse, abuse, or neglect
- Damage caused by parts that were not manufactured or sold by HP
- Damage caused when warranted parts were repaired or replaced by an organization other than HP or by a service provider not authorized by HP

## Software Product Limited Warranty

The HP Software Product Limited Warranty will apply to all Software that is provided to you by HP as part of the HP SureStore E Disk Array FC60 for the NINETY (90) day period specified below. This HP Software Product Limited Warranty will supersede any non-HP software warranty terms that may be found in any documentation or other materials contained in the computer product packaging with respect to covered Software.

**Ninety-Day Limited Software Warranty.** HP warrants for a period of NINETY (90) DAYS from the date of the purchase that the Software will execute its programming instructions when all files are properly installed. HP does not warrant that the software will be uninterrupted or error free. In the event that this software product fails to execute its programming instructions during the warranty period, Customer's remedy will be a refund or replacement. Should HP be unable to replace the media within a reasonable amount of time, Customer's alternate remedy will be a refund of the purchase price (license fee) upon return of the product and all copies.

**Removable Media (if supplied).** HP warrants the removable media, if supplied, upon which this product is recorded to be free from defects in materials and workmanship under normal use for a period of NINETY (90) DAYS from the date of purchase. In the event the media proves to be defective during the warranty period, Customer's remedy will be to return the media to HP for replacement. Should HP be unable to replace the media with a reasonable amount of time, Customer's alternate remedy will be a refund of the purchase price upon return of the product and destruction of all other non-removable media copies of the software product.

Note that removable media does not include "hot-swap" hard drives, which are warranted under the HP Hardware Limited Warranty.

**Notice of Warranty Claims.** All warranty claims must be made during the applicable warranty period or within thirty (30) days after the expiration of the applicable warranty period. Any claim made after that time will not be eligible for warranty service but may be service under a separate HP support contract covering the product.

## Limitation of Warranty

HP does not warrant uninterrupted or error-free operation of an HP SureStore E Disk Array FC60.

This warranty extends only to the original owner in the original country of purchase and is not transferable. Consumables, such as batteries, have no warranty.

The above warranties will not apply to products from which serial numbers have been removed or to defects resulting from misuse (including operation of HP SureStore E Disk Array FC60 without covers and incorrect input voltage), unauthorized modification, operation or storage outside the environmental specifications for the product, in-transit damage, improper maintenance, or defects resulting from use of third-party software, accessories, media, supplies, consumables, or such items not designed for use with the product.

The HP warranty does not cover errors, malfunctions, or problems caused by or related to third-party products that are external to your HP SureStore E Disk Array FC60. HP reserves the right to charge for services where the cause of the error is found to be a non-certified, third-party product. In addition, HP reserves the right to charge for resolving problems caused by unauthorized personnel who have serviced your computer equipment.

**Limitation of liability and Remedies.** In no event will HP or its affiliates, subcontractors, or suppliers be liable for any of the following:

1. Damages for loss of data, or software restoration;
2. Damages relating to Customer's procurement of substitute products or services (i.e., “cost of cover”);
3. Incidental, special, or consequential damages (including lost profits, loss of use, Customer downtime, cost of data recovery/re-creation), even if HP is informed of their possibility;
4. Third-party claims against Customer for losses or damages.

The HP SureStore E Disk Array FC60 is not specifically designed, manufactured or intended for sale as parts, components or assemblies for the planning, construction, maintenance, or direct operation of a nuclear facility. Customer is solely liable if Products or Support purchased by Customer are used for these applications.

HP’s liability for damage to tangible property per incident is limited to the greater of $300,000 or the actual amount paid for the product that is the subject of the claim, and for damages for bodily injury or death, to the extent that all such damages are determined by a court of competent jurisdiction to have been directly caused by a defective HP product.

**TO THE EXTENT ALLOWED BY LOCAL LAW, THE REMEDIES IN THIS WAR-**

**RANTY STATEMENT ARE CUSTOMER'S SOLE AND EXCLUSIVE REMEDIES. EXCEPT AS INDICATED ABOVE, IN NO EVENT WILL HP OR ITS SUPPLIERS BE LIABLE FOR LOSS OF DATA OR FOR DIRECT, SPECIAL, INCIDENTAL, CONSEQUENTIAL (INCLUDING LOST PROFIT OR DATA), OR OTHER DAMAGE, WHETHER BASED IN CONTRACT, TORT, OR OTHERWISE.**

**Some Countries, states or provinces do not allow limitations on the duration of an implied warranty, so the above limitation or exclusion might not apply to you. The warranty gives you specific legal rights and you might also have other rights that vary from country to country, state to state or province to province.**

**THE WARRANTY TERMS CONTAINED HERE, EXCEPT TO THE EXTENT LAWFULLY PERMITTED, DO NOT EXCLUDE, RESTRICT OR MODIFY AND ARE IN ADDITION TO THE MANDATORY STATUTORY RIGHTS APPLICABLE TO THE SALE OF THIS PRODUCT TO YOU.**



## Hewlett-Packard Software License Terms

The disk array hardware described herein uses Licensed Internal Code ("LIC"). The LIC, including any updates or replacements, any LIC utility software, and Supplier's software are collectively referred to as "Software". HP or its Suppliers own the copyrights in the Software and all copies of the Software. HP distributes the Software under license. The following Terms and Conditions shall govern the Software and Customer's use of the Software.

### 1. DEFINITIONS

a. "Software" is one or more programs capable of operating on a computer, processor, or controller which is either listed separately as a Software Product on the Price List, included with another Product on the Price List, or fixed in hardware and not removable in normal operation.

b. "Use" means storing, loading, installing, executing, or displaying Software on a computer, processor, or controller, or making a copy of Software for archival or backup purposes only.

c. "Products" include hardware, Software, options, documentation, accessories, supplies, spare parts, and upgrades on HP's Price List on the date HP receives Customer's order.

**License Grant**. HP grants you a license to Use one copy of the Software. "Use" means storing, loading, installing, executing or displaying the Software. You may not modify the

Software or disable any licensing or control features of the Software. If the Software is licensed for "concurrent use", you may not allow more than the maximum number of authorized users to Use the Software concurrently. You may not allow the Software to be used by any other party or for the benefit of any other party.

**Ownership.** The Software is owned and copyrighted by HP or its third party suppliers. Your license confers no title or ownership and is not a sale of any rights in the Software, its documentation or the media on which they are recorded or printed. Third party suppliers may protect their rights in the Software in the event of any infringement.

**Copies.** You may only make one copy of the Software for archival purposes or when copying is an essential step in the authorized Use of the Software on a backup product, provided that copies are used in no other manner and provided further that Use on the backup product is discontinued when the original or replacement product becomes operable. You must reproduce all copyright notices in the original Software on all copies. You may not copy the Software onto any public or distributed network.

**No Disassembly or Decryption.** You may not disassemble, decompile or reverse engineer the Software without HP's prior written consent. Where you have other rights under statute, you will provide HP with reasonably detailed information regarding any intended disassembly or decompilation. You may not decrypt the Software unless necessary for the legitimate use of the Software.

**Transfer**. You may not sell, sublicense or otherwise transfer the Software. Your license will automatically terminate upon any attempt to transfer of the Software. If you transfer, sell or otherwise dispose of the product upon which this Software is fixed, you must erase the Software before any such transfer or disposal. You will make reasonable efforts to protect the confidentiality of the Software.

**Termination**. HP may terminate your license upon notice for failure to comply with any of these License Terms. Upon termination or when you cease to use the Software, you must immediately destroy the Software, together with all copies and merged portions in any form.

### Restricted Rights Legend

1.) Use, duplication or disclosure by the U.S. Government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause of DFARS 252.227-7013.

**Hewlett-Packard Company**
**3000 Hanover Street**
**Palo Alto, Ca 94304 U.S.A.**

**Copyright © 1997, 1998 Hewlett-Packard Company. All Rights Reserved.**

2.) Customer further agree that Software is delivered and Licensed as "Commercial Computer Software" as defined in DFARS 252.227-7013, or "restricted computer software" as defined in FAR 52.227-19 c(1,2) if used, respectively, in the performance of a Department of Defense on non-Department of Defense U.S. Government contract.

c. Any term of this Agreement which is held to be invalid will be deleted, but the remainder of the Agreement will not be affected.

d. Neither party may assign any rights or obligations without prior written consent of the other party.

e. Customer who exports Products assumes responsibility for complying with applicable laws and regulations and for obtaining required export and import authorization. Customer will not export or re-export the Software, any copy, or any technical data in violation of any applicable laws or regulations.

f. Any disputes arising in connection with this agreement will be governed by the laws of California.

g. These HP Software License Terms supersede any previous communications, representations or agreements between the parties, whether oral or written, regarding transactions hereunder.

# Regulatory Compliance

## Safety Certifications:

- UL listed
- CUL certified
- TUV certified with GS mark
- Gost Certified
- CE-Mark

## EMC Compliance

- US FCC, Class A
- CSA, Class A
- VCC1, Class A
- BCIQ, Class A
- CE-Mark
- C-Tick Mark

# FCC Statements (USA Only)

The Federal Communications Commission (in 47 CFR 15.105) has specified that the following notice be brought to the attention of the users of this product.

This equipment has been tested and found to comply with the limits for a Class A digital device, pursuant to Part 15 of the FCC Rules. These limits are designed to provide reasonable protection against harmful interference when the equipment is operated in a commercial environment. This equipment generates, uses, and can radiate radio frequency energy and, if not installed and used in accordance with the instruction manual, may cause harmful interference to radio communications. Operation of this equipment in a residential area is likely to cause harmful interference in which case the user will be required to correct the interference at his own expense. The end user of this product should be aware that any changes or modifications made to this equipment without the approval of the manufacturer could result in the product not meeting the Class A limits, in which case the FCC could void the user's authority to operate the equipment.



## IEC Statement (Worldwide)

This is a CISPR 22 Class A product. In a domestic environment, this product may cause radio interference, in which case the user may be required to take adequate measures.

## CSA Statement (For Canada Only)

This Class A digital apparatus meets all requirements of the Canadian Interference-Causing Equipment Regulations.

Cet appareil numérique de la classe A respecte toutes les exigences du Règlement.

## VCCI Statement (Japan)

この装置は、情報処理装置等電波障害自主規制協議会（ＶＣＣＩ）の基準に基づくクラスＡ情報技術装置です。この装置を家庭環境で使用すると電波妨害を引き起こすことがあります。この場合には使用者が適切な対策を講ずるよう要求されることがあります。

This equipment is in the Class A category information technology equipment based on the rules of Voluntary Control Council For Interference by Information Technology Equipment (VCCI). When used in a residential area, radio interference may be caused. In this case, user may be required to take appropriate

## Harmonics Conformance (Japan)

高調波ガイドライン適合品

## Class A Warning Statement (Taiwan)

警告使用者：這是甲類的資訊產品，在居住的環境中使用時，可能會造成射頻干擾，在這種情況下，使用者會被要求採取某些適當的對策。

## Spécification ATI Classe A (France Seulement)

DECLARATION D'INSTALLATION ET DE MISE EN EXPLOITATION d'un matèriel de traitement de l'information (ATI), classé A en fonction des niveaux de perturbations radioélectriques émis, définis dans la norme européenne EN 55022 concernant la Compatibilité Electromagnétique.

Cher Client,

Conformément à la Réglementation Française en vigueur l'installation ou le transfert d'installation, et l'exploitation de cet appareil de classe A, doivent faire l'objet d'une déclaration (en deux exemplaires) simultanément auprès des services suivants:

- Comité de Coordination des Télécommunications 20, avenue de Ségur-75700 PARIS
- Préfecture du département du lieu d'exploitation

Le formulaire à utiliser est disponible auprès des préfectures.
La déclaration doit être faite dans les 30 jours suivant la mise en exploitation.
Le non respect de cette obligation peut être sanctionné par les peines prévues au code des Postes et Télvcommunications et celles indiquées dans la loi du 31 mai 1993 susvisée.
Arrêté du 27 Mars 1993, publié au J.O. du 28 Mars-ATI

## Product Noise Declaration (For Germany Only)

**Acoustic Noise Emissions:**

- LpA: 45.0 dB (seeking)
- At bystander position per ISO 7779.
- All data are the results from type tests of the product configuration having the highest acoustic emissions: 12 disk modules.
- All other configurations have lower emission levels.

## Geräuschemission (For Germany Only)

- LpA: 45.0 dB (suchend)
- Am fiktiven Arbeitsplatz nach DIN 45635 T. 19.
- Die Daten sind die Ergebnisse von Typprüfungen an Gerätekonfigurationen mit den höchsten Geräuschemissionen:12 Plattenlaufwerke.
- Alle andere Konfigurationen haben geringere Geräuschpegel.
- Für weitere Angaben siehe unter Umgebungsbedingungen.

# Declaration of Conformity

**according to ISO / IEC Guide 22 and EN 45014**

**Manufacturer Name:** Hewlett-Packard Company
**Manufacturer Address:** Enterprise Storage Business Unit
P.O.Box 15
Boise, Idaho U.S.A. 83707

**Declares, that the product**

**Product Name:** SureStore E Disk Array FC60
**Product Numbers:** A5277A, A5277AZ, A5635A
**Product Options:** All

conforms to the following Product Specifications:

**Safety:** IEC 950: 1991 +A1+ A2 +A3 + A4 / EN 60950 (1992) + A1 + A2 + A3 + A4 + A11
IEC 825-1: 1993 / EN 60825-1: 1994 Class 1 (Laser/LED) GB4943 (1995)

**EMC:** CISPR 22: 1993 / EN 55022 (1994) Class A
EN 50082-1: 1992 - Generic Immunity, including:
IEC 801-2: 1991: / prEN 55024-2 (1991): 4 kV CD, 8 kV AD
IEC 801-3: 1984: / prEN 55024-3 (1991): 3 V/m
IEC 801-4: 1988: / prEN 55024-4 (1992): 1 kV Power Lines, 0.5 kV Data Line
IEC 801-5: 1992: / IEC 61000-4-5 (1995): 2 kV OC, 1 kV SC (Common Mode), 1 kV Peak (Differential Mode)
IEC 61000-3-2: 1997 / EN 61000-3-2 (1995): Harmonic Current
IEC 61000-3-3: 1994 / EN 61000-3-3 (1995): Voltage Fluctuations and Flicker
IEC 61000-4-6: 1996 / EN 61000-4-6 (1996): Conducted Immunity
GB9254 (1988)

**Supplementary Information**

The product herewith complies with the requirements of the Low Voltage Directive 73/23/EEC and the EMC Directive 89/336/EEC and carries the CE Marking accordingly.
1.)This product was tested with Hewlett-Packard Unix server host computer system.



Boise Idaho U.S.A., 03/22/99 Dan T. Michauld / QA Manager

**European Contact: Your local Hewlett-Packard Sales and Service office or Hewlett-Packard GmbH, Department HQ-TRE, Herrenberger Straße 130, D-71034 Boblingen (FAX +49-7031-14-3143)**

# GLOSSARY

**adapter**

A printed circuit assembly that transmits user data (I/Os) between the host system's internal bus and the external Fibre Channel link and vice versa. Also called an I/O adapter, FC adapter, or host bus adapter (HBA).

**ArrayID**

The value used to identify a disk array when using Array Manager 60. The ArrayID can be either the disk array S/N, or an alias assigned to the disk array.

**auto failover**

The process of transferring the I/O data path from a failed controller to the remaining operational controller in the disk array. Auto failover allows uninterrupted access to all user data with no downtime. The failed controller can be removed, then replaced by a new array controller, without any downtime, loss of data, or interruption to the host computer system.

**battery backup unit (BBU)**

A hardware component providing up to 5 days (120 Hours) of backup power to the disk array should power be lost. The BBU power maintains the integrity of the contents of write cache, which is written to the disk media when power to the disk array is restored. A functioning BBU is required to enable disk array cache.

**BCC module**

The BCC (bus controller card) module manages the disk modules in the disk systems. Each disk system contains two BCC controllers.

**bind**

The process of configuring unassigned disks into a LUN disk group. Disks can be bound into one of the following LUN disk groups: RAID 5, RAID 1 (single mirrored pair), RAID 0/1 (multiple mirrored pairs).

**bootware**

This controller firmware comprises the bring-up or boot code, the kernel or executive under which the firmware executes, the firmware to run hardware diagnostics, initialize the hardware and to upload other controller firmware/software from Flash memory, and the XMODEM download functionality.

**bus**

A physical connection between printed-circuit boards in a CPU or a disk array used to communicate data. See also *SCSI-2 bus*.

**caching**

The technique of temporarily storing data in RAM on the disk array controller to improve performance during an I/O. By using higher-speed RAM whenever possible, I/Os can be executed faster than accessing the disk.

**cascaded FC-AL hubs**

One FC-AL hub connected to another FC-AL hub to increase arbitrated loop distances. Cascaded hubs allow distances up to 500 meters between hubs, or between a hub and a device.

**channel**

One of the six SCSI busses connecting the disk array controller enclosure to the disk enclosures. Each channel is an independent LVD SCSI bus.

**channel:ID**

The two-part value used to identify a disk module within the disk array when using the array management tools. The channel value indicates the SCSI channel that the disk enclosure containing the disk is connected to. The ID value indicates the SCSI ID assigned to the slot in which the disk is installed.

**Class of Service**

The types of services provided by the Fibre Channel topology and used by the communicating port.

**controller**

A removable unit that contains an array controller.

**dacstore**

A region on each disk used to store configuration information. During the Start Of Day process, this information is used to configure controller NVSRAM and to establish other operating parameters, such as the current LUN configuration.

**data redundancy**

The technique of storing extra data along with the original data to protect against a disk failure. This extra, or redundant, data is used to rebuild the original data if a disk fails or becomes inaccessible. Data redundancy is achieved using one of two techniques: mirroring or parity.

**data transfer rate**

Data transfer rate is the speed at which data is transferred between a host computer system and a peripheral.

**degraded mode**

The operating mode a LUN enters when if a disk fails or becomes inaccessible. In degraded mode the LUN must recreate the data on the missing disk using either mirrored data or parity information. I/O performance typically suffers while the LUN is in degraded mode. In addition, the LUN is susceptible to a second disk failure while in degraded mode. The LUN will remain in degraded mode until the content of the failed disk is rebuilt to a global hot spare or a replacement disk.

**DIMM**

Dual-inline memory module.

**disk array controller**

A printed-circuit board with memory modules that manages the overall operation of the disk array. The disk array controllers manage all aspects of disk array operation, including I/O transfers, data recovery in the event of a failure, and management of disk array capacity. There are two controllers (A and B) in the disk array enclosure. Both controllers are active, each assuming ownership of LUNs within the disk array.

**disk group**

A collection of individual disks that share a common role in disk array operation. All disks on the disk array will be a member of one of the following disk groups: a LUN, hot spare, or unassigned.

**disk module**

The individual disk drive assemblies used in the disk array. Each module houses a single disk. Up to ten disk modules can be installed in each disk system.

**disk striping**

A performance-enhancing technique in which data is distributed in uniformly sized segments across all the disks in a LUN. Striping allows multiple sets of read/write heads to simultaneously execute an I/O, resulting in improved performance. RAID 5 LUNs use disk striping.

**disk stripe size**

In a LUN using disk striping, the stripe size is calculated by multiplying the stripe segment size times the number of disks in the LUN.

**disk system**

The enclosure that contains the disk modules used by the disk array. Each disk system includes its own BCC controllers, power supplies, and cooling fans. Up to six disk systems can be connected to the disk array.

**Fibre Host ID**

Numbers that identify each controller to the host or other devices on the Fibre Channel network.

**EPROM**

Erasable Programmable Read-Only Memory.

**fabric**

A Fibre Channel term that describes a crosspoint switched network, which is one of three existing Fibre Channel topologies. A fabric consists of one or more fabric elements, which are switches responsible for frame routing. A fabric can interconnect a maximum of 244 devices. The fabric structure is transparent to the devices connected to it and relieves them from responsibility for station management.

**FC-AL**

See *Fibre Channel Arbitrated Loop (FC-AL).*

**fibre**

A generic Fibre Channel term used to cover all transmission media specified in the Fibre Channel Physical Layer standard (FC-PH), including optical fibre, copper twisted pair, and copper coaxial cable.

**fibre optic cable**

An optical fibre cable made from thin strands of dielectric material such as glass through which data in the form of light pulses is transmitted by laser or LED. Fibre optic cable is used for high-speed transmission over medium to long distances.

**Fibre Channel**

Logically, a bidirectional, full-duplex, point-to-point, serial data channel structured for high performance capability. Physically, Fibre Channel interconnects devices such as host systems and servers, FC hubs, and disk arrays through ports, called N_Ports, in one of three topologies: a point-to-point link, an arbitrated loop, or a crosspoint switched network, which is called a fabric. Fibre Channel can interconnect two devices in a point-to-point topology, from two to 126 devices in an arbitrated loop, and up to 224 devices in a fabric switched topology. The disk array default topology is arbitrated loop.

Fibre Channel is a generalized transport mechanism that has no protocol or native I/O command set, but can transport any existing protocol, such as SCSI in Fibre Channel frames. Fibre Channel is capable of operating at speeds of 100 MB/s (full speed), 50 MB/s (half speed),

25 MB/s (quarter speed), or 12.5 MB/s (eighth speed) over distances of up to 100 m over copper media, or up to 10 km over optical links. The disk array operates at full speed.

**Fibre Channel Arbitrated Loop (FC-AL)**

One of three existing Fibre Channel topologies in which two to 126 ports are interconnected serially in a single loop circuit. Access to the FC-AL is controlled by an arbitration scheme. The FC-AL topology supports all classes of service and guarantees in-order delivery of Fibre Channel frames when the originator and responder are on the same FC-AL. The disk array's default topology is arbitrated loop.

**Fibre Channel Arbitrated Loop Hub**

A full-duplex, 1.063 Gb/s intelligent hub used in an FC-AL topology to increase the loop's reliability, the number of loop connections, and the distances between the host system(s) and disk array(s). A maximum of ten devices can be connected to each FC-AL hub.

**Fibre Channel Protocol for SCSI (FCP)**

FCP defines a high-level Fibre Channel mapping layer (FC-4) using lower-level Fibre Channel (FC-PH) services to transmit SCSI command, data, and status information between a SCSI initiator and a SCSI target across the Fibre Channel link using Fibre Channel frame and sequence formats.

**flushing cache**

The process of writing the contents of write cache memory to a disk. Flushing occurs at regular intervals during normal operation, and when power to the disk array is interrupted.

**foreign controller**

Any controller other than the last one to complete the Start-Of-Day process in a given slot. At the completion of the SOD process, the identification of the controller is stored. If a different controller is installed in the slot, it is considered foreign. See *resident controller*.

**frame**

The smallest indivisible unit of application-data transfer used by Fibre Channel. Frame size depends on the hardware implementation and is independent of the application software. Frames begin with a 4-byte Start of Frame (SOF), end with a 4-byte End of Frame (EOF), include a 24-byte frame header and 4-byte Cyclic Redundancy Check (CRC), and can carry a variable data payload from 0 to 2112 bytes, the first 64 of which can be used for optional headers.

**FRU (Field Replaceable Unit)**

A disk array hardware component that can be removed and replaced by a customer or Hewlett-Packard service representative.

**global hot spare**

A disk that is powered up and electrically connected to a disk array but not used until a disk failure occurs. A global hot spare is a dedicated, on-line, backup disk that cannot be used to store user data during normal disk array operation. If any disk in a LUN fails, the disk array automatically begins rebuilding the failed disk's content on an available global hot spare. When the rebuild process completes, the LUN functions normally, using the global hot spare as a replacement for the failed disk. Up to six disks (one for each channel) can be assigned as global hot spares.

**hardware path**

See *primary disk array path*.

**HBA**

See *adapter*.

**high availability**

The technique of designing systems that remain operational even if a hardware or other type of failure occurs. The disk array achieves high availability by using redundant data and global hot spare disks to maintain access to user data in the event of a disk failure. Redundant hardware assemblies also ensure high availability in the event of a failure.

**host**

A processor that runs an operating system using a disk array for data storage and retrieval.

**hot swappable**

Hot swappable components can be removed and replaced while the disk array is online without disrupting system operation. Disk array controller modules, disk modules, power supply modules, and fan modules are all hot swappable components.

**I/O operation**

An operation initiated by a host computer system during which data is either written to or read from a peripheral.

**image (disk image)**

See *mirroring.*

**JBOD**

An acronym for Just a Bunch Of Disks, it implies an enclosure that houses disk drives that by themselves do not provide data redundancy. The disk systems on the disk array are JBOD devices, but when connected to the disk array controllers, they provide data redundancy through the use of RAID technology.

**link**

In Fibre Channel, two unidirectional fibres transmitting in opposite directions and their associated transmitters and receivers that serve as the communication media between nodes in a topology. A link is comparable to a bus in the SCSI protocol.

**logical unit (LUN)**

An acronym for Logical Unit Number, the SCSI term for logical disk unit. A logical unit is an addressable part of a SCSI target. The terms “logical unit number” and “LUN” are interchangeable.

The LUN is the basic storage entity created on the disk array. Each LUN appears to the host as a single hard disk drive. Multiple LUNs, each using different RAID levels, can be

created on the same disk array. A numeric value is assigned to a LUN at the time it is created.

**LVD-SCSI**

Low voltage differential implementation of SCSI. Also referred to as Ultra2 SCSI.

**LVM (Logical Volume Manager)**

The default disk configuration strategy on HP-UX. In LVM one or more physical disk modules are configured into volume groups that are then configured into logical volumes.

**loop address**

The unique ID of a node in Fibre Channel loop topology, sometimes referred to as a Loop ID.

**Major Event Log**

A more detailed logging method implemented on HP08 and later firmware. Major event logging requires that the UTM be enabled.

**MIA**

Media interface adapter, a device that allows you to connect fiber optic cable to a copper wire connector.

**mirrored disks**

A mirrored disk is an exact copy of another disk. Mirrored disk are used when implementing the data redundancy technique of mirroring. See *mirroring*.

**mirroring**

A technique which creates data redundancy by maintaining a duplicate copy of all data stored on the disk array. The mirror copy of data is available if the disk containing the original data fails or becomes inaccessible. RAID 1 and RAID 0/1 LUNs use mirroring.

**NVSRAM**

The disk array controller stores operating configuration information in this non-volatile SRAM (referred to as NVSRAM). The contents of NVSRAM can only be accessed or changed using special diagnostic tools.

**path**

See *primary disk array path or primary path*.

**parity**

A data protection technique that provides data redundancy by creating extra data based on the original data. Parity is calculated on each write I/O by doing a serial binary exclusive OR (XOR) of the data segments in the stripe written to the data disks in the LUN. Parity is used by RAID 5 LUNs to reconstruct data from a failed disk.

**peripheral device addressing (PDA)**

The addressing technique used by the host to address the disk array controllers. Because it is limited to addressing only eight LUNs, PDA has been superseded by VSA as the technique for addressing disk array LUNs.

**point-to-point**

One of three existing Fibre Channel topologies, in which two devices are directly connected by a link with no fabric, loop, or switching elements present.

**port**

The hardware entity that connects a device to a Fibre Channel topology. A device can contain one or more ports.

**primary disk array path** or **primary path**

The main data path used for host I/Os to a LUN. The primary path can be set using LVM. Because the disk array has two controllers, one controller acts as the primary path, and the other acts as the alternate path. If a failure occurs in the primary data path, the alternate path is used, maintaining access to the LUN.

**PROM (Programmable Read-Only Memory)**

SP-resident boot code that loads the SP microcode from one of the disk array's database drives when the disk array is powered up or when an SP is enabled.

**RAID**

An acronym for "Redundant Array of Independent Disks." RAID was developed to provide data redundancy using independent disk drives. RAID is essentially a method for configuring multiple disks into a logical entity (LUN) that appears to the host system as a single, contiguous disk drive. RAID uses features such data striping, disk mirroring, and parity to implement data redundancy. The RAID levels supported by the disk array include *RAID 1, RAID 3,* and *RAID 5*.

**RAID 0**

A RAID level that provides improved performance through the use of data striping, but does not provide any data redundancy. RAID 0 should only be used for non-critcal data storage.

**RAID 1**

A RAID level in which the LUN uses a single mirrored pair of disks. One disk serves as the data disk, and the other serves as the mirror disk.

**RAID 0/1**

A RAID configuration in which the LUN uses both mirroring for redundancy, and disk striping for performance. Half the disks serve as the data disks, and half serve as the mirror disks. Up to 30 disks can be included in a RAID 0/1 LUN.

**RAID 5**

A RAID configuration in which the LUN uses data parity to for redundancy, and disk striping for performance. RAID 5 is efficient in its use of disk space, but suffers a performance penalty when performing write I/Os. Up to six disks (one per enclosure) can be used in a RAID 5 LUN.

**rebuild**

The process of reconstructing the data that was on a failed disk onto another disk, typically a global hot spare. The disk array reconstructs the data using mirrored data or

parity information, depending on the RAID level of the LUN. Until a rebuild is complete, the disk array is operating in degraded mode and is vulnerable to a second disk failure.

**reconstruction**

See *rebuild*.

**resident controller**

The last controller to complete the Start-Of-Day process in a given slot. At the completion of the SOD process, the identification of the controller is stored. This controller remains the resident controller until another controller completes the SOD process in the same slot. See *foreign controller*.

**SAM (System Administration Manager)**

A host-based system configuration and management utility.

**SCSI**

An acronym for “Small Computer System Interface”, SCSI is an industry-standard protocol for connecting peripherals and hosts over a bus topology.

**SCSI-2 bus**

A bus that complies with the SCSI standard. The six channel that connect the disk array controller system to the disk systems are SCSI 2 busses. Each disk system has two internal SCSI-2 busses that can be configured as independent busses that connect five disk slots to each BCC controller. or as a single internal bus.

**SCSI ID**

A unique number assigned to each device connected to a SCSI bus. This number is used by the *HBA* to address each device on the bus. Each disk in a disk system is assigned a SCSI ID of 0 - 9.

**secondary disk array path** or **secondary path**

See *primary disk array path or primary path*.

**SIMM (Single In-line Memory Module)**

A memory module that provides the local storage (cache) for an SP. An SP must have at least two 4-MB memory modules to support the storage system cache.

**Start Of Day (SOD)**

The initialization process used by the disk array controllers to configure itself and establish various operating parameters. Each controller runs its own SOD process. The SOD process occurs following a power on reset, or following the insertion of a controller.

**status light**

Lights on the front panel of each disk module, power module, and fan used to indicate the operating status of the hardware.

**STM (Support Tools Manager)**

A host-based tool used for disk array configuration, management, and diagnostics.

**stripe boundary crossing**

In disk striping, if the stripe segment size is too small for a single I/O operation, the operation requires access to two stripes. Called a stripe boundary crossing, this event reduces I/O performance.

**stripe segment size**

The amount of information simultaneously read from or written to each disk in a LUN using disk striping. The default stripe segment size is set to the same value used for cache page size (4K or 16K). The stripe segment size is configurable and can affect I/O performance. If the segment size for the LUN is set too small, it may result in many stripe boundary crossings.

**striping (disk striping)**

See *disk striping*.

**terminator**

An electrical connection at each end of the SCSI bus composed of a set of resistors (or possibly other components). Its function is to provide a pull-up for open-collector

drivers on the bus, and also impedance matching to prevent signal reflections at the ends of the cable. The SCSI bus requires termination at both ends of the bus. One end of the SCSI bus is terminated by the adapter's internal termination. The other end should have a terminator placed on the 68-pin high density SCSI connector on the last SCSI peripheral. If this device is not terminated, data errors may occur.

**topology**

The physical layout of devices on a network. The three Fibre Channel topologies are fabric, arbitrated loop, and point-to-point. The disk array's default topology is arbitrated loop.

**unbind**

Unbinding reverses the LUN binding process, deleting the LUN and returning the disk in the LUN to the unassigned disk group. Unbinding destroys all data on the LUN.

**Uninterruptible Power Supply (UPS)**

An Uninterruptible Power Supply is a power supply that is capable of maintaining power even if the input ac mains supply loses its source of power.

**Universal Transport Mechanism (UTM)**

A special LUN reserved for communication between the host and the disk array. SCSI commands are passed through the UTM. By default LUN 31 is used for the UTM. If the UTM is disabled, LUN 0 is used for the SCSI command path.

**VHDCI**

Very high density cable interface.

**volume set addressing (VSA)**

An enhanced technique for addressing disk array LUNs. VSA overcomes the eight LUN limit imposed by PDA, allowing all 32 LUNs on the disk array to be addressed by the host.

# INDEX


**A**

A5628A, Array Manager 60 20
AC power specifications 425
  disk enclosure 421
acoustics
  controller enclosure 423
  disk enclosure 427
adapter
  Fibre Channel host 411
adding a disk module from another array 255
adding a global hot spare
  using Array Manager 60 296
  using SAM 273
  using STM 315
adding capacity 254
  with disk enclosures 224, 256
  with disk modules 254
  with higher capacity disk modules 256
adding disk enclosures 256
adding disk modules 254
address
  Fibre Channel host ID 179
addressing
  array modules 208
  full-bus example 194
  LUNs 208
  peripheral device addressing (PDA) 208
  split-bus example 194
  volume set addressing (VSA) 208
alias, disk array
  assigning using Array Manager 60 297
  assigning using SAM 264
alternate data path 64
alternate LUN path
  definition 64
AM60Srvr
  starting 241
AM60Srvr daemon 241
amcfg
  binding a LUN 289
  changing LUN ownership 293
  unbinding a LUN 292
amdsp
  checking rebuild progress 305
  listing disk arrays 292
amlog
  viewing logs 309
ammgr
  adding hot spare 296
  assigning an alias 297
  displaying parity scan status 307
  halting parity scan 307
  performing a parity scan 306
  removing hot spare 296
  resetting battery age 312
  setting cache flush limit 301
  setting cache flush threshold 300
  setting cache page size 300
  setting controller date and time 297
amutil
  changing rebuild priority settings 305
  flushing disk array log 311
  locating disk modules 304
  managing the UTM 298
AR60Diag
  using to create a boot device 222
Array Manager 60 20
  adding a global hot spare 296
  assigning an alias 297
  binding a LUN 289

calculating LUN capacity 292
changing LUN ownership 293
changing rebuild priority settings 305
checking disk array status 282
checking rebuild progress 305
command summary 278
described 238
displaying parity scan status 307
flushing disk array log 311
halting parity scan 307
identifying disk modules 292
installing 240, 241
listing all disk arrays 288
locating disk modules 304
managing disk array logs 309
managing log files 307
performing a parity scan 306
removing a global hot spare 296
replacing a LUN 294
rescanning for disk arrays 288
resetting battery age 312
running 241
setting cache flush limit 301
setting cache flush threshold 300
setting cache page size 300
synchronizing controller date and time 297
unbinding a LUN 293
array *See disk array*
ArrayID 280
listing for all arrays 288
assigning LUN ownership 247

**B**

backplane controller card. *See BCC module*
basic topology 102, 103
error recovery 108
batteries *See battery backup module*
battery age
resetting 312
battery backup module
cache protection 46
described 45
life expectancy 46
removal and replacement 400
battery charger *See battery backup module*
BCC module
described 29
troubleshooting 380
binding a LUN
using Array Manager 60 289
using SAM 267
using STM 314
boot support 222

**C**

cabinets
supported 18
cable installation
Fibre Channel 196
SCSI 187
cables
FC-AL supported lengths 110
length
performance degradation 116
cache
upgrading to 512 MB 258
cache flush limit
impact on performance 252
setting 301
cache flush threshold
impact on performance 251
setting 300
cache page size
impact on performance 252
setting 300
caching 66
read I/O 66
write I/O 66
calculating LUN capacity 292
campus topology 102, 125
capacity
adding 254
expanding with disk enclosures 224
managing 242

maximum 75
changing LUN ownership 293
channel number
disk module 244
channel:ID
described 280
checking disk array status
using Array Manager 60 282
using SAM 260
using STM 314
circuit breakers, European 149
cntrlrID
described 280
command summary
Array Manager 60 278
command-based interface *See cstm*
component status conditions 321
configuration
management software 216
switches 176
configuring LUNs 242
controller date
synchronizing with host 297
controller enclosure
acoustics 423
dimensions 420
environmental 422
front cover 37
host ID 39
installation 170, 173
LED locations 204
models 412
power down 43
power supply fan module
described 43
power supply modules
described 42
specifications 420
voltage 421
weight 420
controller enclosure modules
described 34
controller fan module
described 40
removal and replacement 398
controller memory modules
DIMM 40
SIMM 40
controller module
described 38
interface connectors 39
LEDs, described 39
controller time
synchronizing with host 297
cstm
described 351
current
disk enclosure 425
inrush 147
steady state 147
total operating and in-rush 148

**D**

data channel, verifying 206
data parity
described 48
data striping
described 49
device logs
STM 351
dimensions
controller enclosure 420
disk enclosure 424
DIMM 40
disabling disk WCE 302
disk array
capacity management features 65
configuration planning 71
described 18
high availability features 47
high availability planning 73
management software 213
options 414
power input specifications 425

power-down sequence 205
power-up sequence 198
rebuild process 61
upgrade and add-on products 416
using as a boot device 222
ventilation 403
disk array capacity
maximum 75
disk array configurations
five disk enclosure, high availability and performance 92
five disk enclosure, maximum capacity 94
four disk enclosure, high availability and performance 88
four disk enclosure, maximum capacity 90
one disk enclosure, non-high availability 78
recommended 77
six disk enclosure, high availability and performance 96
six disk enclosure, maximum capacity 98
three disk enclosure, high availability and performance 84
three disk enclosure, maximum capacity 86
two disk enclosure, high availability and performance 80
two disk enclosure, maximum capacity 82
disk array logs
managing 309
viewing 309
disk array management software 20
disk array status
checking using Array Manager 60 282
checking using SAM 260
checking using STM 314
disk drives *See disk modules*
disk enclosure
acoustics 427
adding 256
current 425
described 23
dimensions 424
heat output 425
hot swappable modules 27
installation 160
moving 168
operating with empty slot 28
operation features 25
options 414
power 425
power input specifications 421
power on switch 199
specifications 424
status indicators 25
status indicators, described 25
voltage 425
weight 424
disk enclosure modules
described 27
disk filler
described 27
disk mirroring
described 47
disk module
adding 67, 254, 386
adding or replacing 386
channel number 244
described 27
enclosure number 245
identifying in array 244
identifying using Array Manager 60 292
installation 166
moving 255
removal and replacement 386
replacing with higher capacity modules 256
SCSI ID 244
slot number 245
disk *See disk module*
disk status conditions 319
disk system *See disk enclosure*
disk WCE
disabling 302
enabling 303
disks
alternate LUN path 64

primary LUN path 64
drive lockout 387
drivers
system 146

**E**
electrical requirements 147
EMC compliance 434
EMS hardware event monitoring 21, 362
enabling disk WCE 303
enclosure number 245
environmental requirements
disk enclosure electrical requirements 149
electrical 147
power distribution units (PDUs/ PDRUs) 150
recommended European circuit breakers 149
recommended PDU/PDRU for HP System/E racks 151
site 147
environmental specifications 422, 426
ESD strap
part number 160
evaluating performance 250
event messages, contents of 363
event notification
on HP-UX 363
expanding storage capacity
adding disk enclosures 224
overview 75
Expert Tool 355
menu options 356
using menu mode 356
using X Windows 355

**F**
factory default configurations *See recommended array configurations*
fan
troubleshooting 381
fan module
troubleshooting 380
fan module, disk enclosure
described 31
fault detection 31
removal and replacement 392
fast write cache LED 40
FC-AL, *See* Fibre Channel Arbitrated Loop
Fibre Channel
controller modules 39
host connections 39
hub connection 71
IO adapters and drivers 146
throughput 71
Fibre Channel Arbitrated Loop 102, 131
Fibre Channel cable
installation 196
Fibre Channel connection 71
Fibre Channel host ID
addressing 179
filler module
adding or replacing 386
installation 166
removal
Firmware Update Tool 357
firmware, BCC
troubleshooting 382
flushing disk array log 311
front cover, controller enclosure 37
removal and replacement 397
FRU codes 322
full-bus mode
setting 178
switch 178

**G**
global hot spare
adding using Array Manager 60 296
adding using SAM 273
adding using STM 315
removing using Array Manager 60 296
removing using SAM 275
removing using STM 315

tips for selecting disks 62
global hot spare disks
described 61

**H**
hardware event monitoring *See EMS hardware event monitoring*
hardware path
interpreting 208
peripheral device addressing 208
sample ioscan 207
volume set addressing (VSA) 209
heat output
controller enclosure 421
disk enclosure 425
high availability
features 21, 47
planning 73
high availability topology 102
error recovery 118
hardware components 115
redundant HP FC-AL Hubs 115
high availability, distance, and capacity
topology 102, 120–123
error recovery 123
host
connections 39
host adapters 146
host ID
controller modules 39
hot spare *See global hot spare*
hot swappable modules 18
disk enclosure 27
HP Fibre Channel Arbitrated Loop Hub 102
campus topology 125
cascading 120
high availability topology 115
high availability, distance and capacity topology 120
single-system distance topology 110
HP-UX boot device 222

**I**
I/Os per second
performance 74
identifying disk modules 244
using Array Manager 60 292
increasing capacity
of a LUN 245
of the disk array 254
information log
STM 354
Information Tool 352
interpreting information log 354
using menu mode 353
using X windows 352
in-rush current 147
installation
controller enclosure 170
disk enclosures 160
disk enclosures into rack 162
disk module 166
filler mdoule 166
installing Array Manager 60 software 240
interpreting status indicators 264
ioscan output 207
ioscan verification 206

**L**
LEDs
controller enclosure normal operation 203, 369
controller module, described 39
disk enclosure normal operation 201
disk enclosures, described 25
fast write cache 40
functions 377
power supply module 395
status 379
listing all disk arrays 288
locating disk modules
using Array Manager 60 304
using SAM 265
using STM 265, 316

log files
managing 307
logs
managing 309
loop ID *See Fibre Channel host ID*
losing LUN 0 376
LUN
addressing 208
assigning ownership 247
binding using Array Manager 60 289
binding using SAM 267
binding using STM 314
calculating capacity using Array Manager 60 292
changing ownership using Array Manager 60 293
configuring 242
described 65
replacing using Array Manager 60 294
selecting disks for 243
selecting RAID level for 247
setting stripe segment size 249
unbinding using Array Manager 60 293
unbinding using SAM 271
unbinding using STM 315
LUN 0, losing 376
LUN numbers
determining from VSA hardware path 211
LUN status conditions 317
LVM
data paths 64

**M**
management software
disk array 213
system requirements 213
management tools 20, 238
managing disk array capacity 242
managing LUNs
using Array Manager 60 289
using SAM 267
using STM 314
menu-based interface *See mstm*
MIA
installation 197
mirroring
described 47
missing LUN 0 376
models and options
controller enclosure 412
modules
controller enclosure 34
disk enclosure 27
moving a disk enclosure 168
moving disk modules
from one array to another 255
within an array 255
mstm
described 349

**N**
noise level 423

**O**
operating environment 426
specifications 422
operating system
support 410
verification 214
operating system support 20
overheating
controller enclosure 40

**P**
parity
described 48
parity scan 306
displaying status of 307
halting 307
path, verification 206
PDA (peripheral device addressing) 208
PDU/PDRU 150
installation 152
order numbers 417

recommended for HP System/E racks 151
troubleshooting 380
performance
array configuration 73
I/Os per second 74
impact of configuration settings 250
rebuild 61
SCSI channels 72
split-bus configurartion 74
peripheral device addressing (PDA) 208
planning a disk array configuration 71
power
AC input specifications 425
AC input, disk enclosure 421
DC specifications 422, 426
disk enclosure 425
recommended UPS 152
power cable
troubleshooting 380
power consumption
controller enclosure 421
disk enclosure 425
power cord
installation 183
power cord, controller enclosure
with power supply 42
power distribution units (PDU/PDRU) 150
power down
controller enclosure 43
power supply
troubleshooting 381
power supply fan module, controller enclosure
described 43
removal and replacement 403
power supply module
LEDs 395
power supply module, controller enclosure
described 42
removal and replacement 405
power supply module, disk enclosure
described 32
load sharing 32
removal and replacement 394
power switch
controller enclosure 200
disk enclosure 199
disk enclosure, location 26
power, controller enclosure
redundancy, controller enclosure 42
power-down sequence 205
powering up the disk array 198
power-up
controller 366
LEDs 366
self tests 366
primary data path 64
primary LUN path
definition 64
products
add-on & support 416
proper ventilation 403

**R**

rack
Rittal 18
racks
supported 18
RAID
described 47
supported levels 21
RAID 0
described 50
RAID 0/1
described 52
five disk enclosure array 92
four disk enclosure array 88
six disk enclosure array 96
RAID 1
described 51
three disk enclosure array 84
two disk enclosure array 80
RAID 3
described 53
RAID 5

described 55
five disk enclosure array 94
four disk enclosure array 90
six disk enclosure array 98
storage capacity 75
three disk enclosure array 86
two disk enclosure array 82
RAID level comparison
application and I/O pattern performance characteristics 60
data redundancy characteristics 57
general performance characteristics 59
RAID levels
described 50
tips on selecting 247
rail installation
HP original rack 162
rebuild 61
changing priority settings 305
checking progress of 305
recommended disk array configurations 77
recommended UPS 152
regulatory compliance 434
removal and replacement
battery backup module 400
controller fan module 398
fan module, disk enclosure 392
front cover, controller enclosure 397
power supply fan module, controller enclosure 403
power supply module, controller enclosure 405
power supply module, disk enclosure 394
SCSI cables 407
removing a global hot spare
using Array Manager 60 296
using SAM 275
using STM 315
replacing a LUN
using Array Manager 60 294
using SAM 271
replacing disk modules with higher capacity modules 256
rescanning for disk arrays 288
Rittal rack 18
running Array Manager 60 241

**S**

safety compliance 434
SAM 260
adding a global hot spare 273
assigning an alias 264
binding a LUN 267
checking disk array status 260
interpreting status indicators 264
locating disk modules 265
removing a global hot spare 275
replacing a LUN 271
unbinding a LUN 271
scanning for parity errors 306
SCSI
drive connections 39
SCSI cables 187
SCSI cabling
full-bus configurations 188
removal and replacement 407
split-bus configurations 191
SCSI channels
number of connections 72
performance 72
throughput 72
SCSI connections 178
SCSI ID
disk module 244
SCSI sense codes 327
segment size
impact on performance 250
tips on setting 249
selecting a RAID level 247
selecting disks for a LUN 243
Sense Code Information 317
sense codes, SCSI 327
setting
cache flush limit 301

cache flush threshold 300
cache page size 300
configuration switches 176
controller date and time 297
stripe segment size 249
SF21 384
SF88 384
SIMMs 40
single-system distance topology 102, 110
D-Class, K-Class, T-Class, and V-Class 110
error recovery 113
high availability 111
non-high availability 111
site requirements 147
slot number 245
software 21
installing disk array management 213
requirements 213
management 20
software configuration 216
software tools 20
specifications 424
controller enclosure 420
disk enclosure 424
split-bus mode
performance 74
switch 178
status
checking using Array Manager 60 282
checking using SAM 261
checking using STM 314
status conditions
component 321
disk 319
LUN 317
Status Conditions and Sense Code Information 317
status LEDs, controller enclosure
normal operation 203, 369
status LEDs, disk enclosure
normal operation 201
steady state current 147
STM
adding a global hot spare 315
binding a LUN 314
checking disk array status 314
cstm 351
device logs 351
Expert Tool 355
Firmware Update Tool 357
Information Tool 352
locating disk modules 265, 316
managing with 314
mstm 349
removing a global hot spare 315
status indications 379
unbinding a LUN 315
users interfaces 347
xstm 348
storage capacity
expanding with disk enclosures 224
expansion 75
scalable 21
storage capacity options 414
storage expansion
adding disks 67
stripe segment size
tips on setting 249
supported operating systems 20
SureStore E Disk System SC10 *See disk enclosure*
switch configuration 176
disk enclosure 187
disk enclosure (tray) ID switch 176
Fibre Channel host ID address 179
full-bus 178
split-bus 178
tips on setting 176
switch configurations, Fibre Channel 129
synchronizing
controller date and time 297
system requirements
specifications 410

**T**
throughput
Fibre Channel 71
SCSI channels 72
topologies
unsupported Windows 131
topology
basic 102, 103
error recovery 108
campus 102, 125
high availability 102
error recovery 118
hardware components 115
redundant HP FC-AL Hubs 115, 118
high availability, distance, and capacity 102, 120–123
error recovery 123
single-system distance 102, 110, 113
error recovery 113
high availability 111
non-high availability 111
switch configurations 129
troubleshooting
checklist 365
event notification 363
isolating causes 380
power-up 366
table 380
Windows NT tools 362

**U**
unbinding a LUN
using Array Manager 60 293
using SAM 271
using STM 315
universal transport mechanism 298
unsupported topologies
Windows 131
upgrade and add-on products 416
upgrading controller cache 258
UPS, recommended 152
user interfaces for STM 347
UTM 298

**V**
ventilation 403
verifying host to array path 206
viewing logs 309
voltage specifications
controller enclosure 421
disk enclosure 425
volumn set addressing (VSA) 208

**W**
warranty information 428
WCE
disabling on disks 302
disk 302, 303
enabling on disks 303
weight
controller enclosure 420
disk enclosure 424
Windows NT tools 20
Windows NT troubleshooting tools 362
write cache
battery backup 46

**X**
X Windows interface *See xstm*
xstm
described 348